The peer-reviewed version of this paper is published in Simulation Modelling Practice and Theory (http://dx.doi.org/10.1016/j.simpat.2016.12.013). This version is typeset by the authors and differs only in pagination and typographical detail.

arXiv:1509.09174v4 [cs.OH] 6 Jan 2017

Model-independent comparison of simulation output

Nuno Fachada[1], Vitor V. Lopes[2], Rui C. Martins[3], and Agostinho C. Rosa[1]

[1] *Institute for Systems and Robotics, LARSyS, Instituto Superior Técnico, Universidade de Lisboa, Lisboa, Portugal*

[2] *UTEC - Universidad de Ingeniería & Tecnología, Lima, Jr. Medrano Silva 165, Barranco, Lima, Perú*

[3] *INESC TEC, Campus da FEUP, Rua Dr. Roberto Frias, 4200-465 Porto, Portugal*

Abstract

Computational models of complex systems are usually elaborate and sensitive to implementation details, characteristics which often affect their verification and validation. Model replication is a possible solution to this issue. It avoids biases associated with the language or toolkit used to develop the original model, not only promoting its verification and validation, but also fostering the credibility of the underlying conceptual model. However, different model implementations must be compared to assess their equivalence. The problem is, given two or more implementations of a stochastic model, how to prove that they display similar behavior? In this paper, we present a model comparison technique, which uses principal component analysis to convert simulation output into a set of linearly uncorrelated statistical measures, analyzable in a consistent, model-independent fashion. It is appropriate for ascertaining distributional equivalence of a model replication with its original implementation. Besides model-independence, this technique has three other desirable properties: a) it automatically selects output features that best explain implementation differences; b) it does not depend on the distributional properties of simulation output; and, c) it simplifies the modelers' work, as it can be used directly on simulation outputs. The proposed technique is shown to produce similar results to the manual or empirical selection of output features when applied to a well-studied reference model.

Keywords— Model alignment; Docking; PCA; Model replication; Simulation output analysis

1 Introduction

Complex systems are usually described as consisting of mutually interacting objects, often exhibiting complex global behavior resulting from the interactions between these objects. This behavior is typically characterized as "emergent" or "self-organizing" as the system's constituting parts do not usually obey a central controller [1]. Analytic treatment generally does not yield the complete theory of a complex system. Therefore, modeling and simulation techniques play a major role in our understanding of how these systems work [2]. Methodologies such as agent-based modeling (ABM), system dynamics, discrete event simulation, among others, are frequently employed for this purpose [3]. Of these, ABM provides an instinctive approach for describing many complex systems, as agents are regularly a suitable match to the individual and heterogeneous objects composing these systems. The local interactions of these objects, as well as their individual and adaptive behavior, are often critical for understanding global system response [4, 5]. ABMs are commonly implemented as a stochastic process, and thus require multiple runs (observations) with distinct pseudo-random number generator (PRNG) seeds in order to have appropriate sample sizes for testing hypotheses and differentiating multiple scenarios under distinct parameterizations [6].

Computational models of complex systems in general, and ABMs in particular, are usually very sensitive to implementation details, and the influence that seemingly negligible aspects such as data structures, discrete time representation and sequences of events can have on simulation results is

notable [7]. Furthermore, most model implementations are considerably elaborate, making them prone to programming errors [8]. This can seriously affect model validation[1] when data from the system being modeled cannot be obtained easily, cheaply or at all. Model verification[2] can also be compromised, to the point that wrong conclusions may be drawn from simulation results.

A possible answer to this problem is the independent replication of such models [8]. Replication consists in the reimplementation of an existing model and the replication of its results [11]. Replicating a model in a new context will sidestep the biases associated with the language or toolkit used to develop the original model, bringing to light dissimilarities between the conceptual and implemented models, as well as inconsistencies in the conceptual model specification [12, 9]. Additionally, replication promotes model verification, model validation [9], and model credibility [11]. More specifically, model verification is promoted because if two or more distinct implementations of a conceptual model yield statistically equivalent results, it is more likely that the implemented models correctly describe the conceptual model [9]. Thus, it is reasonable to assume that a computational model is untrustworthy until it has been successfully replicated [12, 13].

Model parallelization is a an illustrative example of the importance of replication. Parallelization is often required for simulating large models in practical time frames, as in the case of ABMs reflecting systems with large number of individual entities [14]. By definition, model parallelization implies a number of changes, or even full reimplementation, of the original model. Extra care should be taken in order to make sure a parallelized model faithfully reproduces the behavior of the original serial model. There are inclusively reports of failure in converting a serial model into a parallel one [15].

Although replication is considered the scientific gold standard against which scientific claims are evaluated [16], most conceptual models have only been implemented by the original developer, and thus, have never been replicated [17, 9, 8, 11]. Several reasons for this problem have been identified, namely: a) lack of incentive [8, 11]; b) below par communication of original models [16, 18]; c) insufficient knowledge and uncertainty of how to validate results of a reimplemented model [9]; and, d) the inherent difficulty in reimplementing a model [12, 9, 8]. This work targets c), with positive influence on d). Replication is evaluated by comparing the output of the reimplementation against the output of the original model [11], and this process, as will be discussed, is empirically driven and model-dependent (or even parameter-dependent). Furthermore, it is sometimes unclear as to what output features best describe model behavior. A robust and ready to use output comparison method would thus reduce or eliminate uncertainty of how to validate reimplementation results (reason c), eliminating this obstacle in the overall process of model replication (reason d).

We present a model comparison technique, which uses principal component analysis (PCA) [19] to convert simulation output into a set of linearly uncorrelated statistical measures, analyzable in a consistent, model-independent fashion. It is appropriate for ascertaining statistical equivalence of a model replication with its original implementation. Besides model-independence, this technique has three additional desirable features: a) it automatically selects output features that best explain implementation differences; b) it does not depend on the distributional properties of simulation output; and, c) it simplifies the modelers' work, as it can be used directly on model output, avoiding manual selection of specific points or summary statistics. The proposed method is presented within the broader context of comparing and estimating the statistical equivalence of two or more model replications using hypothesis tests. The technique is evaluated against the empirical selection of output summary statistics, using the PPHPC ABM [20] as a test case. This model is thoroughly studied in terms of simulation output for a number of parameter configurations, providing a solid base for this discussion.

The paper is organized as follows. First, in Section 2, we review commonly used methods for comparing the output of simulation models, as well as previous work on model replication using these methods. An overview of hypothesis testing within the scope of model output comparison is conducted in Section 3. The main steps in designing and performing a model comparison experiment using hypothesis tests are presented in Section 4. The proposed model-independent comparison methodology

[1] Determining if the model implementation adequately represents the system being modeled [9] for its intended purpose [10].

[2] Determining if the model implementation corresponds to a specific conceptual model [9].

is described in Section 5. Section 6 introduces PPHPC, the simulation model used as a test case for the proposed model comparison approach. Section 7 delineates the experimental setup for assessing this methodology against the manual selection of output summary measures. In Section 8, results of the empirical and model-independent comparison approaches are presented. A discussion and an evaluation of these results is performed in Section 9. Recommendations on using the proposed method are given in Section 10. Section 11 summarizes what was accomplished in this paper.

2 Background

Axtell et al. [17] defined three replication standards (RSs) for the level of similarity between model outputs (Carley [21] calls the RS the emphasis of demonstration): *numerical identity*, *distributional equivalence* and *relational alignment*. The first, *numerical identity*, implies exact numerical output, but it is difficult to demonstrate for stochastic models and not critical for showing that two such models have the same dynamic behavior. To achieve this goal, *distributional equivalence* is a more appropriate choice, as it aims to reveal the statistical similarity between two outputs. Finally, *relational alignment* between two outputs exists if they show qualitatively similar dependencies with input data, which is frequently the only way to compare a model with another which is inaccessible (e.g., implementation has not been made available by the original author), or with a non-controllable "real" system (such as a model of the human immune system [22]). For the remainder of this text we assume the *distributional equivalence* RS when discussing model replication.

For the *distributional equivalence* RS, a set of statistical summaries representative of each output are selected. It is these statistical summaries, and not the complete outputs, that will be compared in order to assert the similarity between the original model and the replication. As models may produce large amounts of data, the summary measures should be chosen as to be relevant to the actual modeling objective. The summaries of all model outputs constitute the set of focal measures (FMs) of a model [9], or more specifically, of a model parameterization (since different FMs may be selected for distinct parameterizations). There are three major statistical approaches used to compare FMs: 1) statistical hypothesis tests; 2) confidence intervals; and, 3) graphical methods [23].

Statistical hypothesis tests are often used for comparing two or more model implementations [17, 9, 12, 24, 25, 14]. More specifically, hypothesis tests check if the statistical summaries obtained from the outputs of two (or more) model implementations are drawn from the same distribution. However, statistical testing may not be the best option for comparing the output of a model with the output of the system being modeled. Since the former is only an approximation of the latter, the null hypothesis that FMs from both are drawn from the same distribution will almost always be rejected. Although statistical procedures for comparing model and system outputs using hypothesis tests have been proposed [26], confidence intervals are usually preferred for such comparisons, as they provide an indication of the magnitude by which the statistic of interest differs from model to system. Confidence intervals are also commonly used when evaluating different models that might represent competing system designs or alternative operating policies [23, 27]. Graphical methods, such as Q–Q plots, can also be employed for comparing output data, though their interpretation is more subjective than the previous methods.

A number of simulation models have been replicated and/or aligned by using statistical methods to compare FMs. In reference [17], Axtell et al. compared two initially different models, with one iteratively modified in order to be aligned with the other. The authors evaluated how can distinct equivalence standards be statistically assessed using non-parametric statistical tests (namely the Kolmogorov–Smirnov [28] and Mann–Whitney U [29] tests), and how minor variations in model design affect simulation outcomes. They concluded that comparing models developed by different researchers and with different tools (i.e., programming languages and/or modeling environments), can lead to exposing bugs, misinterpretations in model specification, and implicit assumptions in toolkit implementations. The concepts and methods of "computational model alignment" (or "docking") were first discussed in this work.

Edmonds and Hales [12] performed two independent replications of a previously published model involving co-operation between self-interested agents. Several shortcomings were found in the original model, leading the authors to conclude that unreplicated simulation models and their results cannot be trusted. This work is one of the main references in model replication, describing in detail the process of running two model implementations with different parameters, selecting comparison measures and performing adequate statistical tests.

In reference [9], the authors presented an ABM replication case study, describing the difficulties that emerged from performing the replication and determining if the replication was successful. A standard t-test was used for comparing model outputs. The authors concluded that model replication influences model verification and validation and promotes shared comprehension concerning modeling decisions.

Miodownik et al. [24] replicated an ABM of social interaction [30], originally implemented in MATLAB, using the PS-I environment for ABM simulations [31]. A statistical comparison of the mean levels of "civicness" at the end of the simulation (over 10 runs) was performed using the Mann–Whitney U test. Results showed that, while distributional equivalence was obtained in some cases, the two models were mostly only relationally aligned. The authors attribute this mainly to the fact that some aspects of the original model were not implementable with PS-I.

A method for replicating insufficiently described ABMs was discussed in reference [25], consisting in modeling ambiguous assumptions as binary parameters and systematically applying statistical tests to all combinations for their equivalence to the original model. The approach was used to replicate Epstein's demographic prisoner's dilemma model [32], with only partial success, suggesting the existence of some undefined assumptions concerning the original model. The authors also conducted a number of statistical tests regarding the influence of specific design choices, highlighting the importance that these be explicitly documented.

Alberts et al. [33] implemented a CUDA [34] version of Toy Infection Model [35], and compared it with the original version implemented in NetLogo [36], as well as to an in-house optimized serial version. Statistical validation was performed visually using Q–Q plots.

Multiple serial and parallel variants of the PPHPC model were compared in reference [14]. Simultaneous comparison of the several variants was accomplished by applying the multi-group non-parametric Kruskal–Wallis test [37] to predetermined FMs of the several model outputs. Results showed that all model variants could not be considered misaligned for a number of different parameters.

3 Statistical hypothesis tests for comparing FMs

As described in the previous section, hypothesis tests are commonly used for comparing FMs collected from two or more model implementations. Since statistical tests are a central subject matter in this discussion, the current section aims to clarify their use in the context of model comparison. More specifically, this section illustrates: a) how to generically interpret statistical tests; b) what tests to use and how their choice strongly depends on the statistical properties of FMs; c) what techniques are available to assess these properties; and, d) how to deal with the case of multiple FMs and/or test results.

In statistical hypothesis testing, a null hypothesis (H_0) is tested against an alternative hypothesis (H_1). The test can have two results: 1) fail to reject H_0; or, 2) reject H_0 in favor of H_1. In the most simple terms, the result of a statistical test is derived from the p-value it generates. The p-value is the probability of obtaining a result equal or more unexpected than what was observed in the data if H_0 is true. If the p-value is below a predefined significance level α, the test result is deemed statistically significant and the null hypothesis, H_0, is rejected. However, the test result may not be clear, especially when the p-value is close to the typically used significance levels, $\alpha = 0.01$ or $\alpha = 0.05$. Thus, it may be preferable to show the p-value instead of presenting the test result as a binary decision.

The rejection of a true H_0 is designated as a false positive or a type I error. Consequently, it stems from the p-value definition that the significance level, α, also represents the type I error rate. A type II

error, or false negative, occurs when the test fails to reject a false H_0. The type II error rate, denoted as β, has an inverse relation with α, such that there is a trade-off between the two. The statistical power of a test, defined as $1 - \beta$, is the probability that it correctly rejects a false H_0. Reducing the type I error rate, α, increases β, and consequently reduces the power of a test. Increasing the sample size is a general way to increase the power of a test without modifying α [38].

In the context of model comparison, the tests of interest are two-sample (or multi-sample) statistical tests which test for the null hypothesis that the observations in each sample are drawn from the same distribution, against the alternative that they are not. Here, samples or groups correspond to the model implementations being compared. Typically, when more than one FM is to be compared, univariate tests are used to compare individual FMs. Nonetheless, FMs may also be combined into one multidimensional random variable, and then simultaneously compared using a multivariate test. In either case, the choice of test also depends on characteristics of the FM and on the assumptions that can be made concerning its underlying distribution. Tests which make such assumptions are labeled as parametric tests. Such tests are generally more powerful than their non-parametric counterparts, which do not expect any particular underlying distribution. Table 1 lists tests which, according to these aspects, are commonly used to check if samples are drawn from the same distribution.

	Parametric	Non-parametric
Univariate ($s = 2$ groups)	t-test [38]	Mann–Whitney [29] Kolmogorov–Smirnov [28]
Univariate ($s > 2$ groups)	ANOVA [38]	Kruskal–Wallis [37]
Multivariate	MANOVA [39, 40]	Various [41, 42, 43, 44, 45]

Table 1 – Hypothesis tests commonly used to check if samples are drawn from the same distribution. Tests are organized by parametric assumptions, sample dimensionality and number of groups (samples) which can be compared.

The tests listed in Table 1 expect that samples are mutually independent. This can be guaranteed if samples are properly collected, as discussed later in Section 4.2. The parametric tests make more stringent assumptions, though, also requiring that [46, 47, 39]: a) each sample is drawn from a normally distributed population (multivariate normality for MANOVA); and, b) samples are drawn from populations with equal variances (for MANOVA, the variance–covariance matrix should be the same for each population). These assumptions can be checked using additional statistical tests. More specifically, group sample normality can be assessed using the Shapiro–Wilk test [48] (Royston test [49] for multivariate samples), while equality of variances among groups can be verified with the Bartlett test [50] (Box's M test [51] for homogeneity of variance–covariance matrices in the multivariate case). However, Box's M test is very sensitive and can lead to false negatives (type II errors). Fortunately, MANOVA is considerably robust to violations in the homogeneity of variance–covariance matrices when groups have equal sample sizes [40].

If these assumptions are not met, non-parametric tests may be used instead. Non-parametric tests may be also preferable if: a) a specific FM is better represented by the median (the listed parametric tests check for differences in means); b) the sample size is small; c) the data is not continuous; or, d) the data contains outliers.

If the choice falls on multivariate tests there are a few caveats. For example, in MANOVA each FM is a dependent variable (DV). However, MANOVA is not appropriate when: a) DVs are highly correlated, which may be the case for FMs derived from outputs of the same simulation run; and, b) when the number of DVs or dimensions is higher than the number of observations (i.e., when there are more FMs than runs per model implementation). Additionally, the non-parametric multivariate alternatives (e.g., [41, 42, 43, 44, 45]) are not as widespread and well established as MANOVA, and

are commonly oriented towards specific research topics.

When more than one FM is to be compared and the choice of test falls on univariate tests, several p-values will be generated, as one test will be applied per FM. This is known as a multiple comparisons problem. Since the p-value is the probability of obtaining a result equal or more unexpected than what was observed if H_0 is true, the generation of multiple p-values increases the likelihood of false positives, i.e., of rejecting a true H_0 in some of the performed tests. For example, with $\alpha = 0.05$, it is expected that approximately 5% of the tests incorrectly reject H_0. This problem can be addressed with a multiple testing correction procedure, such as the Bonferroni or Holm methods [52]. However, multiple comparison correction methods often assume independency of test statistics, which might not be possible to assure when testing different outputs of the same simulation model, which are most likely correlated. Additionally, such approaches may increase the likelihood of type II errors. As such, it is often preferable to simply present the p-values and discuss possible interpretations, such that the reader can draw his own conclusions [53].

In the case of models with multiple outputs and multiple statistical summaries per output, there may exist an intermediate logical grouping of FMs. More specifically, FMs can be logically grouped by output. In such cases, does the researcher ignore this logical grouping of FMs, and bundles them together in a multivariate or multiple univariate comparison problem? Or does the researcher treat FMs from different outputs separately? This a model-dependent issue, and generally, the simplest course of action is to perform a univariate test per FM, present all the (uncorrected) p-values, and discuss the results considering the increased likelihood of type I errors.

4 Designing and performing a model comparison experiment

In order to setup a model comparison experiment it is necessary to first define a research question, which can be stated as follows:

> Determine if two or more implementations of a conceptual simulation model display statistically indistinguishable behavior.

If two or more implementations produce the same dynamic behavior it is reasonable to expect that any output collected over a number of independent runs also comes from the same distribution [12]. Thus, the research question can be reformulated in the following manner:

> Determine if two or more implementations of a conceptual simulation model generate statistically indistinguishable outputs, given the same inputs.

As described in Section 2, output similarity is assessed using statistical summaries representative of each output, i.e., by verifying if the selected FMs are distributionally equivalent. Furthermore, to answer the original question with a higher degree of confidence, output similarity should be observed for distinct input parameter sets [12]. Consequently, the research question can be further refined:

> Determine if two or more implementations of a conceptual simulation model generate statistically indistinguishable FMs for several parameterizations.

Procedure 1 summarizes the process of performing a model comparison experiment using hypothesis tests to answer this research question. In the following subsections each step of the procedure is analyzed in detail.

4.1 Choose the parameter sets with which to generate FMs

In order to demonstrate that two or more model implementations are misaligned, it is sufficient to show that at least one FMs is statistically different for one parameterization. On the other hand, demonstrating that FMs are drawn from the same distribution for any number of parameter sets does

Procedure 1 Steps for performing a model comparison experiment.

1. Choose the parameter sets with which to generate FMs.
2. For each parameter set:
   - 2.1 Perform n replications for each model implementation and collect the respective outputs.
   - 2.2 Select FMs to compare.
   - 2.3 Extract FMs from collected outputs.
   - 2.4 Analyze statistical properties of FMs.
   - 2.5 Compare FMs using hypothesis tests.
   - 2.6 Decide on the alignment or otherwise of the compared model implementations for the current parameter set.
3. Decide on the global alignment or otherwise of the compared model implementations.

not definitively prove that the compared implementations are aligned. Nonetheless, demonstrating alignment for more parameter sets increases the confidence that the implementations are in fact globally aligned [12].

As a starting point, a minimum of two parameterizations should be used to verify if implementations are aligned. The chosen parameter sets should induce noticeably distinct simulation behaviors, thus triggering different model mechanisms and increasing the probability of finding implementation differences.

4.2 Perform a number of replications for each model implementation and collect the respective outputs

Statistical best practices should be followed when setting up the replications or runs for each model implementation. More specifically: a) there should be enough replicates in order to guarantee adequate statistical power; and, b) the replicates should be independently generated.

Considering the first issue, an appropriate minimum sample size (i.e., number of replicates or runs per model implementation) is required so that the test adequately rejects false null hypotheses. Long run times in more complex models may limit the number of observations which can be collected in practice. While small sample sizes reduce the likelihood that statistically significant results indicate a true difference, large samples can lead to the detection of small and/or irrelevant differences [38]. The determination of sample size for a desired power level depends on the distributional properties of each FM, which in practice could mean that different FMs warrant distinct samples sizes. Additionally, in a two-sample or $s > 2$ sample testing scenario, it is preferable to have equal numbers of observations from the populations being compared, since this may facilitate interpretation of results and make parametric tests more robust to variance homogeneity assumptions [40, 38]. As such, sample size should be chosen by balancing the desired statistical power, the available computational resources, as well as the simplicity and/or convenience of the analysis to be performed. A thorough discussion on determining the minimum sample size for a desired power level in the context of simulation models is presented in reference [6].

The second issue concerns the independence of replicates. This can be achieved by performing each run with a distinct pseudo-random number stream. Typically, the same PRNG algorithm with different seeds is used for this purpose. However, some care is required in order to avoid correlated or overlapped streams. For example, using the run number as the seed should be avoided. A simple solution consists of using a random spacing approach to split the PRNG stream, for example by

applying a cryptographic hash (distinct from the PRNG algorithm used in the simulations) to the run number, and use the resulting value as the PRNG seed [20, 14]. Allying this technique with a long-period PRNG such as the Mersenne Twister [54] minimizes the likelihood of correlated or overlapped streams being used in different runs.

4.3 Select FMs to compare

Output summary measures, or more generally, FMs, are commonly selected in an empirical fashion, and normally consist of long-term or steady-state means. Nonetheless, being limited to analyze average system behavior can lead to incorrect conclusions [27, 6]. Consequently, other measures such as proportions or extreme values can also be used to assess model behavior.

For steady-state FMs, care must be taken with initialization bias, which may cause substantial overestimation or underestimation of the long-term performance [55]. Such problems can be avoided by discarding data obtained during the initial transient period, before the system reaches steady-state conditions. The simplest way of achieving this is to use a fixed truncation point, l, for all runs with the same initial conditions, selected such that: a) it systematically occurs after the transient state; and, b) it is associated with a round and clear value, which is easier to communicate [55]. A moving average procedure for selecting l is discussed in references [20, 27].

4.4 Extract FMs from collected outputs

Let $\mathbf{X}_j = \begin{bmatrix} X_{j0} & X_{j1} & X_{j2} & \dots & X_{jm} \end{bmatrix}$ be an output collected from the j^{th} simulation run of one model implementation. The X_{ji}'s (rows under 'Iterations' in Table 2) are random variables that will, in general, be neither independent nor identically distributed [27], and as such, are not adequate to be used directly in a classical statistical analysis. On the other hand, let $X_{1i}, X_{2i}, \dots, X_{ni}$ be the observations of an output at iteration i for n runs (columns under 'Iterations' in Table 2), where each run begins with the same initial conditions, but uses a different stream of random numbers as a source of stochasticity. The X_{ji}'s will now be independent and identically distributed (IID) random variables, to which classical statistical analysis can be applied. However, because individual values of the output at some iteration i are unlikely to be representative of the output as a whole, summary measures, as shown in Table 2, under 'Statistical summaries', should be used instead.

Rep.	Iterations					Statistical summaries				
1	X_{10}	X_{11}	$\dots$	$X_{1,m-1}$	$X_{1,m}$	$f_1(\mathbf{X}_1)$	$f_2(\mathbf{X}_1)$	$\dots$	$f_{q-1}(\mathbf{X}_1)$	$f_q(\mathbf{X}_1)$
2	X_{20}	X_{21}	$\dots$	$X_{2,m-1}$	$X_{2,m}$	$f_1(\mathbf{X}_2)$	$f_2(\mathbf{X}_2)$	$\dots$	$f_{q-1}(\mathbf{X}_2)$	$f_q(\mathbf{X}_2)$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
n	X_{n0}	X_{n1}	$\dots$	$X_{n,m-1}$	$X_{n,m}$	$f_1(\mathbf{X}_n)$	$f_2(\mathbf{X}_n)$	$\dots$	$f_{q-1}(\mathbf{X}_n)$	$f_q(\mathbf{X}_n)$

Table 2 – Values of a generic simulation output from one model implementation (under 'Iterations') for n replications of m iterations each (plus iteration 0, i.e., the initial state), and the respective summary measures (under 'Statistical summaries'). Values along columns are IID.

The statistical summaries, or more generally, FMs, can be empirically defined as discussed in Section 4.3 or automatically extracted using the method proposed in Section 5. FMs can be considered as functions, $f_1(.)$ to $f_q(.)$, which when applied to individual output observations, return scalar values. Taken column-wise, these values are IID (because they are obtained from IID replications), constituting a *sample* prone to statistical analysis, or more specifically, to hypothesis testing. There will be as many samples (or groups of observations) per FM as there are model implementations. For example, considering some conceptual model with a single output X, FM samples for the h^{th} implementation would be constituted as follows:

$$
\begin{aligned}
\mathbf{f}_1^h &= \begin{bmatrix} f_1(\mathbf{X}_1^h) & f_1(\mathbf{X}_2^h) & \dots & f_1(\mathbf{X}_n^h) \end{bmatrix} \\
\mathbf{f}_2^h &= \begin{bmatrix} f_2(\mathbf{X}_1^h) & f_2(\mathbf{X}_2^h) & \dots & f_2(\mathbf{X}_n^h) \end{bmatrix} \\
&\vdots \\
\mathbf{f}_q^h &= \begin{bmatrix} f_q(\mathbf{X}_1^h) & f_q(\mathbf{X}_2^h) & \dots & f_q(\mathbf{X}_n^h) \end{bmatrix}
\end{aligned}
$$

Alternatively, FMs can be combined into one multivariate FM, e.g., for the h^{th} implementation we would have:

$$
\mathbf{F}^h = \begin{bmatrix} f_1(\mathbf{X}_1^h) & f_2(\mathbf{X}_1^h) & \dots & f_q(\mathbf{X}_1^h) \\ f_1(\mathbf{X}_2^h) & f_2(\mathbf{X}_2^h) & \dots & f_q(\mathbf{X}_2^h) \\ \vdots & \vdots & \ddots & \vdots \\ f_1(\mathbf{X}_n^h) & f_2(\mathbf{X}_n^h) & \dots & f_q(\mathbf{X}_n^h) \end{bmatrix}
$$

where rows correspond to observations, and columns to variables or dimensions.

4.5 Distributional analysis of FMs and choice of statistical tests

In order to choose adequate tests for comparing each FM, it is necessary to first check if the FM samples are in accordance with the assumptions described in Section 3. If an FM appears to follow the normal distribution and if the respective samples have similar variance, then the parametric tests listed in Table 1 will be suitable. Conversely, non-parametric tests will be acceptable if such assumptions are not verified, or if the median is considered to better represent the FM in question.

If the choice falls on a multivariate test, the presentation of results becomes considerably simpler, since there is only one p-value to exhibit. However, it is important to keep in mind that the distributional analysis of the combined multidimensional FM becomes significantly more complex. Thus, it may be difficult to assess the suitability of MANOVA, or even of the non-parametric alternatives, for comparing the multidimensional FM. Furthermore, when a significant difference exists, the multivariate test may mask where such difference occurs, while univariate tests will clearly show which FMs are affected.

4.6 Compare FMs using hypothesis tests

In the univariate case, a statistical test is applied per FM to samples from the compared models. For example, considering two model implementations with three FMs, three tests would be performed. The first test on samples $\mathbf{f}_1^1$ and $\mathbf{f}_1^2$, the second on samples $\mathbf{f}_2^1$ and $\mathbf{f}_2^2$, and the third on samples $\mathbf{f}_3^1$ and $\mathbf{f}_3^2$. Each test returns a p-value, so we could perform a multiple comparisons correction or present the uncorrected p-values keeping in mind the caveat of increased type I error likelihood.

Continuing with the previous example, if FMs are instead combined into a multidimensional FM, then a multivariate test (e.g., MANOVA) would be performed on samples $\mathbf{F}^1$ and $\mathbf{F}^2$, yielding a single p-value.

4.7 Analyze results and decide on the alignment or otherwise of the compared model implementations

Hypothesis tests results should be analyzed per parameterization, because model implementations may be locally aligned for some parameter sets, but not for others. Thus, each parameterization should be considered a separate experiment, i.e., it should not be considered part of a multiple comparisons

scenario. If at least one FM is shown to be statistically different for a given parameterization, then the compared implementations can be considered globally misaligned.

If the p-value(s) for a specific parameterization are above the chosen significance level α, the implementations can be considered locally aligned for that parameter set. When one or more p-values are significant in a multiple comparison scenario (for the same parameterization), they should be in the proportion predicted by α. No p-values should remain significant after Bonferroni-type corrections are performed. In the multivariate approach, where all the FMs are merged into one multidimensional FM, the single p-value should not be significant.

It is important to understand that statistical tests may not provide a definitive answer. Only in the case of clear misalignments, with very significant p-values, can one reject H_0 with confidence. In many situations some p-values may appear to be borderline significant, and is up to the modeler to judge if the detected misalignment has practical significance.

5 Model-independent selection of FMs

The empirical selection of FMs, as described in Section 4.3, has a number of disadvantages. First, it relies on summary measures which are model-dependent and, probably, user-dependent. Furthermore, for different model parameters, the originally selected FMs may be of no use, as simulation output may change substantially. For example, the warm-up period for the steady-state statistical summaries can be quite different. Certain models and/or parameterizations might not even display a steady-state behavior. Finally, it might not be clear which FMs best capture the behavior of a model.

A model-independent approach to FM selection should work directly from simulation output, automatically selecting the features that best explain potential differences between the implementations being compared, thus avoiding the disadvantages of an empirical selection. Additionally, such a method should not depend on the distributional properties of simulation output, and should be directly applicable by modelers.

Our proposal consists of automatically extracting the most relevant information from simulation output using PCA. PCA is a widely used technique [56] which extracts the largest variance components from data by applying singular value decomposition to a mean-centered data matrix. In other words, PCA is able to convert simulation output into a set of linearly uncorrelated measures which can be analyzed in a consistent, model-independent fashion. This technique is especially relevant for stochastic simulation models, as it considers not only equilibrium but also dynamics over time [9]. Procedure 2 summarizes this approach, replacing steps 2.2 and 2.3 of Procedure 1.

Procedure 2 accepts, on a per output basis, the collected data (Procedure 1, step 2.1) in the form of matrix $\mathbf{X}$, yielding matrix $\mathbf{T}$, which contains the representation of $\mathbf{X}_c$ (i.e., the column mean-centered version of $\mathbf{X}$) in the principal components (PCs) space, as well as vector $\boldsymbol{\lambda}$, containing the eigenvalues of the covariance matrix of $\mathbf{X}_c$. The columns of $\mathbf{T}$ correspond to PCs, and are orderer by decreasing variance, i.e., the first column corresponds to the first PC, and so on. Rows of $\mathbf{T}$ correspond to observations. The k^{th} column of $\mathbf{T}$ contains sn model-independent observations for the k^{th} PC, n for each implementation. Thus, each PC corresponds to an FM. As in the case of empirically selected FMs, univariate or multivariate statistical tests can be used to check if samples from different implementations are drawn from the same distribution. However, both testing approaches will not prioritize dimensions, even though the first PCs are more important for characterizing model differences, as they explain more variance. This can be handled in the univariate case by prioritizing PCs according to their explained variance using the weighted Bonferroni procedure on the resulting p-values [57]. For multivariate tests, dimensions/variables can be limited to the number of PCs that explain a prespecified amount of variance, although there is no prioritization within the selected dimensions.

The eigenvalues vector $\boldsymbol{\lambda}$ is also important for this process, for two reasons: 1) to select a number of PCs (i.e., FMs) to be considered for hypothesis testing, such that these explain a prespecified percentage of variance; 2) the alignment or otherwise of s model implementations can be empirically assessed by analyzing how the explained variance is distributed along PCs. The percentage of variance

Procedure 2 Obtaining model-independent summary measures from one generic simulation output from s model implementations.

1. Group the X_{ji}^{h}'s from all replications row-wise in matrix $\mathbf{X}$ for each implementation, as follows:

$$\mathbf{X} = \begin{bmatrix} X_{10}^{1} & X_{11}^{1} & \dots & X_{1,m-1}^{1} & X_{1,m}^{1} \\ X_{20}^{1} & X_{21}^{1} & \dots & X_{2,m-1}^{1} & X_{2,m}^{1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ X_{n0}^{1} & X_{n1}^{1} & \dots & X_{n,m-1}^{1} & X_{n,m}^{1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ X_{10}^{s} & X_{11}^{s} & \dots & X_{1,m-1}^{s} & X_{1,m}^{s} \\ X_{20}^{s} & X_{21}^{s} & \dots & X_{2,m-1}^{s} & X_{2,m}^{s} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ X_{n0}^{s} & X_{n1}^{s} & \dots & X_{n,m-1}^{s} & X_{n,m}^{s} \end{bmatrix}$$

2. Determine matrix $\mathbf{X}_c$, which is the column mean-centered version of $\mathbf{X}$.

3. Apply PCA to matrix $\mathbf{X}_c$, considering that rows (replications) correspond to observations and columns (iterations or time steps) to variables. This yields:

   - Matrix $\mathbf{T}$, containing the representation of the original data in the principal components (PCs) space.

$$\mathbf{T} = \begin{bmatrix} T_{11}^{1} & T_{12}^{1} & \dots & T_{1,u-1}^{1} & T_{1,u}^{1} \\ T_{21}^{1} & T_{22}^{1} & \dots & T_{2,u-1}^{1} & T_{2,u}^{1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ T_{n1}^{1} & T_{n2}^{1} & \dots & T_{n,u-1}^{1} & T_{n,u}^{1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ T_{11}^{s} & T_{12}^{s} & \dots & T_{1,u-1}^{s} & T_{1,u}^{s} \\ T_{21}^{s} & T_{22}^{s} & \dots & T_{2,u-1}^{s} & T_{2,u}^{s} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ T_{n1}^{s} & T_{n2}^{s} & \dots & T_{n,u-1}^{s} & T_{n,u}^{s} \end{bmatrix}$$

   - Vector $\boldsymbol{\lambda}$, containing the eigenvalues of the covariance matrix of $\mathbf{X}_c$ in descending order, each eigenvalue corresponding to the variance of the columns of $\mathbf{T}$.

$$\boldsymbol{\lambda} = \begin{bmatrix} \lambda_1 & \lambda_2 & \dots & \lambda_{u-1} & \lambda_u \end{bmatrix}$$

explained by each PC can be obtained as shown in Eq. 1.

$$S_k^2(\%) = \frac{\lambda_k}{\sum \boldsymbol{\lambda}} \tag{1}$$

where k identifies the k^{th} PC, λ_k is the eigenvalue associated with the k^{th} PC, and $\sum \boldsymbol{\lambda}$ is the sum of all eigenvalues. If the variance is well distributed along many PCs, it is an indication that the compared implementations are aligned, at least for the output being analyzed. On the other hand, if most of the variance is explained by the first PCs, it can be an indication that at least one model implementation is misaligned. The rationale being that if all implementations show the same dynamical behavior, then the projection of their outputs in the PC space will be close together and have similar statistics, i.e., means, medians and variance. As such, PCA will be unable to find components which explain large quantities of variance, and the variance will be well distributed along the PCs. If at least one model implementation is misaligned, the projection of its outputs in the PC space will be farther apart than the projections of the remaining implementations. As such, PCA will yield at least one component which explains large part of the overall variance.

The alignment of two or more implementations can be assessed by analyzing the following information: 1) the p-values produced by the univariate and multivariate statistical tests, which should be above the typical 1% or 5% significance levels in case of implementation alignment; in the univariate case, it may be useful to adjust the p-values using the weighted Bonferroni procedure to account for multiple comparisons; 2) in case of misalignment, the total number of PCs required to explain a pre-specified amount of variance should be lower than in case of alignment; also, more variance should be explained by the first PCs of the former than by the same PCs of the latter; and, 3) the scatter plot of the first two PC dimensions, which can offer visual, although subjective feedback on model alignment; e.g., in case of misalignment, points associated with runs from different implementations should form distinct groups.

5.1 Extension to multiple outputs

Determining the alignment of models with multiple outputs may be less straightforward. If model implementations are clearly aligned or misaligned, conclusions can be drawn by analyzing the information provided by the proposed method applied to each one of the model outputs.

In order to compare all model outputs simultaneously, a Bonferroni or similar p-value correction could be used. This strategy can be directly applicable in the multivariate case, since there is one p-value per model output. However, a direct application to the univariate case would be more complex, since there will be multiple p-values per model output (one p-value per PC), and these may have been previously corrected with the weighted Bonferroni procedure according to their explained variance.

We propose an alternative approach based on the concatenation of all model outputs, centered and scaled. This reduces a model with g outputs to a model with one output, which can be processed with Procedure 2. In order to perform output concatenation, outputs are centered and scaled such that their domains are in the same order of magnitude, replication-wise. This can be performed using range scaling on each output, for example, as shown below for a given simulation output;

$$\widetilde{X}_{ji} = \frac{X_{ji} - \overline{\mathbf{X}}_j}{\max \mathbf{X}_j - \min \mathbf{X}_j}, \qquad i = 0, 1, \ldots, m;\ j = 1, \ldots, sn \tag{2}$$

$$\widetilde{\mathbf{X}}_j = \begin{bmatrix} \widetilde{X}_{j0} & \widetilde{X}_{j1} & \ldots & \widetilde{X}_{jm} \end{bmatrix} \tag{3}$$

where i represents iterations or time steps, j is the replication number, X_{ji} is the output value at iteration i and replication j, $\mathbf{X}_j$ is a row vector with the complete output generated in the j^{th} replication, and $\widetilde{\mathbf{X}}_j$ is its range scaled version. Other centering and scaling methods, such as auto-scaling or level scaling [58], can be used as an alternative to range scaling in Eq. 2. For a model with g outputs, the resulting concatenated output for the j^{th} replication is given by

$$\widetilde{\mathbf{A}}_j = {}_1\widetilde{\mathbf{X}}_j \oplus {}_2\widetilde{\mathbf{X}}_j \oplus \ldots \oplus {}_g\widetilde{\mathbf{X}}_j \tag{4}$$

where $\oplus$ is the concatenation operator, and $\widetilde{\mathbf{A}}_j$ is the concatenation of all model outputs for replication j. Model implementations can thus be compared with the proposed method using the "single" model output.

6 Simulation model

The Predator–Prey for High-Performance Computing (PPHPC) model is a reference model for studying and evaluating spatial ABM (SABM) implementation strategies, capturing important characteristics of SABMs, such as agent movement and local agent interactions. It is used in this work as a test case for the proposed model comparison method. The model is thoroughly described in reference [20] using the ODD protocol [4]. Here we present a summarized description of the model.

6.1 Description

PPHPC is a predator–prey model composed of three entity classes: *agents*, *grid cells* and *environment*. *Agents* can be of type prey or predator. While prey consume passive cell-bound food, predators consume prey. Agents have an energy state variable, E, which increases when feeding and decreases when moving and reproducing. When energy reaches zero, the agent is removed from the simulation. Instances of the *grid cell* entity class are where agents act, namely where they try to feed and reproduce. Grid cells have a fixed grid position and contain only one resource, cell-bound food (*grass*), which can be consumed by prey, and is represented by the countdown state variable C. The C state variable specifies the number of iterations left for the cell-bound food to become available. Food becomes available when $C = 0$, and when a prey consumes it, C is set to c_r (an initial simulation parameter). The set of all grid cells forms the *environment* entity, a toroidal square grid where the simulation takes place. The environment is defined by its size and by the restart parameter, c_r. The temporal scale is represented by a positive integer m, which represents the number of iterations.

Simulations start with an initialization process, where a predetermined number of agents are randomly placed in the simulation environment. Cell-bound food is also initialized at this stage. After initialization, and to get the simulation state at iteration zero, outputs are collected. The scheduler then enters the main simulation loop, where each iteration is sub-divided into four steps: 1) agent movement; 2) food growth in grid cells; 3) agent actions; and, 4) gathering of simulation outputs. Note that the following processes are explicitly random: a) initialization of specific state variables (e.g., initial agent energy); b) agent movement; c) the order in which agents act; and, d) agent reproduction. For process c), this implies that the agent list should be explicitly shuffled before agents can act.

Six outputs are collected at each iteration i: P_i^s, P_i^w, P_i^c, $\overline{E}_i^s$, $\overline{E}_i^w$, and $\overline{C}_i$. P_i^s and P_i^w refer to the total prey (*sheep*) and predator (*wolf*) population counts, respectively, while P_i^c holds the quantity of available cell-bound food. $\overline{E}_i^s$ and $\overline{E}_i^w$ contain the mean energy of prey and predator populations. Finally, $\overline{C}_i$ refers to the mean value of the C state variable in all grid cells.

6.2 Parameterizations

Reference parameters for the PPHPC model are specified in reference [20]. Parameters are qualitatively separated into size-related and dynamics-related groups. Although size-related parameters also influence model dynamics, this separation is useful for parameterizing simulations.

Concerning size-related parameters, a base grid size of 100×100 is associated with 400 prey and 200 predators. Different grid sizes should have proportionally assigned agent population sizes, such that the initial agent density and the initial ratio between prey and predators remains constant. We

define *model size* as the association between grid size and initial agent population. For example, model size 200 corresponds to a grid size of 200×200 with 1600 prey and 800 predators at iteration 0.

For the dynamics-related parameters, two parameter sets, 1 and 2, are proposed. The two parameterizations generate distinct dynamics, with the second set typically yielding more than twice the number of agents than the first during the course of a simulation. We will refer to a combination of model size and parameter set as "size@set", e.g., 400@1 for model size 400, parameter set 1. The reference number of iterations, m, is 4000, not counting with the initial simulation state at iteration 0.

6.3 Empirical selection of FMs

Under the parameterizations discussed in the previous subsection, all outputs have an initial transient stage, entering steady-state after a number of iterations. The point at which simulations enter this state of equilibrium varies considerably from parameter set 1 to parameter set 2, occurring considerably sooner in the former. Model size does not seem to affect the steady-state truncation point, mainly influencing the magnitude of the collected outputs. As such, steady-state is empirically established for $i > 1000$ and $i > 2000$ for parameter sets 1 and 2, respectively [20].

Considering this information, as well as the recommendations discussed in Section 4.3, the following statistical summaries are selected for individual outputs: 1) maximum value (max); 2) iteration where maximum value occurs ($\arg\max$); 3) minimum value (min); 4) iteration where minimum value occurs ($\arg\min$); 5) steady-state mean ($\overline{X}^{\text{ss}}$); and, 6) steady-state sample standard deviation (S^{ss}). Thus, we specify a total of 36 FMs (six statistical summaries from six outputs). Naturally, the steady-state measures are collected after the truncation point defined for each parameter set.

As discussed in reference [20], all measures, except $\arg\max$ and $\arg\min$, are amenable to be compared using parametric methods, as they follow (or approximately follow) normal distributions.

6.4 Implementations

Two implementations of the PPHPC model are used for evaluating the proposed model comparison technique. The first is developed in NetLogo [20], while the second is a Java implementation with several parallel variants [14]. For the results discussed in this paper, simulations performed with the Java implementation were executed with the EX parallelization strategy using eight threads. In this strategy, each thread processes an equal part of the simulation environment, and reproducible simulations are guaranteed. Individual threads use their own sub-sequence of a global random sequence, obtained with a random spacing approach using the SHA-256 cryptographic hash function. The Mersenne Twister pseudo-random number generator (PRNG) [54] is used by both implementations for driving the model's random processes. Complete details of both implementations are available in the provided references, and their source code is available at `https://github.com/fakenmc/pphpc/`.

7 Experimental setup

In order to test the model comparison methods, we define a base PPHPC configuration using the NetLogo implementation and compare it with three configurations using the Java implementation. All configurations are tested for model sizes 100, 200, 400 and 800, and parameter sets 1 and 2, as described in Section 6.2. The four configurations follow the conceptual model, except when stated otherwise:

1. NetLogo implementation.
2. Java implementation.
3. Java implementation: agents are ordered by energy prior to agent actions, and agent list shuffling is not performed.

4. Java implementation: the restart parameter, c_r, is set to one unit less than specified in the tested parameterizations (9 instead of 10 for parameter set 1, 14 instead of 15 for parameter set 2).

The goal is to assess how the two FM selection strategies (empirical and model-independent) expose the increasing differences of comparing configuration 1 with configurations 2–4. More specifically, we are interested in checking if the proposed model-independent comparison method is able to expose these differences in the same way as the manual or empirical approach. As such, we define three comparison cases:

Case I Compare configuration 1 with configuration 2. These configurations should yield distributionally equivalent results.

Case II Compare configuration 1 with configuration 3. A small misalignment is to be expected.

Case III Compare configuration 1 with configuration 4. There should be a mismatch in the outputs.

For each "size@set" combination, independent samples of the six model outputs are obtained from $n = 30$ replications for each configuration, in a total of $4n = 120$ runs. Each replication $r = 1, \ldots, 4n$ is performed with a PRNG seed obtained by taking the MD5 checksum of r and converting the resulting hexadecimal string to an integer (limited to 32-bits for NetLogo and 128-bits for the Java implementation), guaranteeing independence between seeds, and consequently, between replications. The same samples are used for the evaluation of both empirical and model-independent FM selection.

In order to evaluate how the tested methodologies respond to larger sample sizes, an additional $n = 100$ replications were performed per configuration for the 400@1 combination, in a total of $4n = 400$ runs. The PRNG seeds for individual replications were obtained in the same fashion. Similarly, the tested methodologies were also compared with $n = 10$ for the 400@1 combination, but in this case using the first 10 replications from the $n = 30$ setup for each configuration. All samples sizes were chosen for convenience with the purpose of simplifying analysis of results, namely of how the two FM selection approaches, empirical and model-independent, fare under the same conditions.

For the model-dependent comparisons, results were obtained with the `SimOutUtils` MATLAB toolbox [59]. The `micompr` R package [60] provides an implementation of the proposed model-independent comparison approach, and was used to produce the corresponding results.

The data generated by this computational experiment, as well as the scripts used to set up the experiment, are made available to other researchers at `https://zenodo.org/record/46848`.

8 Results

In this section we mainly focus on the results for the 400@1 combination. Results for the remaining size/set combinations are provided as Supplementary material, and are referred to when appropriate.

8.1 Empirical selection of FMs

Table 3 shows the results for the empirical approach to cases I, II and III. The corresponding p-values were obtained with the t-test for max, min, $\overline{X}^{\text{ss}}$ and S^{ss} statistics, and with the Mann–Whitney U test for the arg max and arg min statistics. For the latter, when all points in both samples have the same value, we simply present 1.00 as the p-value, since the test is not performed.

Results for case I show that the two configurations are reasonably well aligned. In a total of 36 FMs, five are below the 5% significance level, and of these, only one is below 1%. Since the p-values are not corrected for multiple comparisons, it is expected that some of them present "significant" differences. For case II, the effect of not shuffling the agents before they act in the Java implementation (configuration 3) is evident, as half of the p-values are below 0.05. Of these, all but one are also below 0.01. Finally, in case III, most p-values show very significant differences, such that we can reject, with

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.213	0.288	0.011	0.231	0.774	0.086
	arg max	0.733	0.802	0.056	0.858	0.284	0.279
	min	0.002	0.088	0.094	0.275	0.076	0.011
	arg min	0.048	0.350	0.279	1.000	0.091	1.000
	$\overline{X}^{ss}$	0.457	0.905	0.546	0.833	0.049	0.551
	S^{ss}	0.242	0.282	0.285	0.474	0.109	0.285
II	max	7e−6	1e−7	2e−4	2e−7	0.322	0.004
	arg max	1e−4	0.540	0.230	0.703	0.052	1e−4
	min	1e−5	4e−4	0.009	0.508	0.017	0.001
	arg min	0.712	0.906	0.001	1.000	0.009	1.000
	$\overline{X}^{ss}$	0.206	< 1e−8	6e−8	< 1e−8	0.938	6e−8
	S^{ss}	0.735	0.324	0.712	0.218	0.688	0.713
III	max	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8
	arg max	2e−5	0.115	< 1e−8	0.134	0.332	1e−5
	min	7e−8	< 1e−8	< 1e−8	0.582	0.003	< 1e−8
	arg min	2e−5	0.070	2e−5	1.000	0.004	< 1e−8
	$\overline{X}^{ss}$	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8
	S^{ss}	3e−7	7e−8	5e−5	4e−8	0.013	0.013

Table 3 – P-values for the empirical selection of FMs for model size 400, parameter set 1, and $n = 30$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{ss}$ and S^{ss} statistics, and with the Mann-Whitney U test for arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

a high degree of confidence, the hypothesis that the corresponding FMs are produced by equivalent configurations.

Tables S1.1–S1.8, provided as Supplementary material, show the results for the remaining size/set combinations. The general tendency is the same, i.e., p-values for case I highlight very few to no significant differences between the tested configurations (although there are, sporadicly, some significant differences in the $\overline{E}^s$ output), while for cases II and III, the dissimilarities are apparent. However, perhaps unexpectedly, results for parameter set 2 indicate that the configurations compared in case II have more significant differences than those in case III.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	24	32	25	36	43	26	39
	MNV	0.528	0.258	0.548	0.105	0.746	0.577	0.704
	t-test	0.530	0.836	0.804	0.784	0.378	0.805	0.879
	PCS	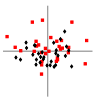	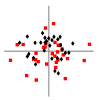	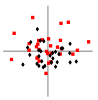	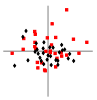	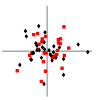	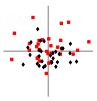	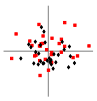
II	#PCs	24	29	24	12	43	25	38
	MNV	≤1e-08	≤1e-08	5e-08	≤1e-08	0.467	1e-07	≤1e-08
	t-test	0.042	≤1e-08	0.108	≤1e-08	0.017	0.109	0.390
	PCS	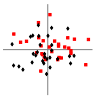	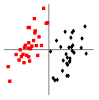	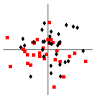		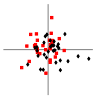	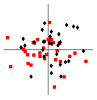	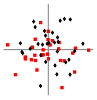
III	#PCs	17	7	8	13	38	1	31
	MNV	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	*	≤1e-08
	t-test	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	PCS	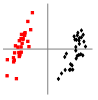						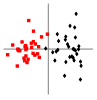

Table 4 – Model-independent comparison for model size 400, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) hypothesis tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replaced with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

8.2 Model-independent selection of FMs

Table 4 presents the results for the model-independent comparison approach for 400@1. Regarding the number of PCs, there is not much difference between case I and II, except for the mean prey energy, $\overline{E}^s$, as this output is clearly the most affected by the introduced change (no agent list shuffling). For case III, however, less PCs are required to explain the same percentage of variance in all outputs, meaning that configurations 1 and 4 are less aligned than configurations 1 and 2 (case I) and configurations 1 and

3 (case II). Nonetheless, the p-values offer more objective answers, displaying increasingly significant differences from case I to case III, i.e., in line with the results from the empirical approach. P-values were obtained with the t-test for the first PC, and with MANOVA for the number of PCs which explain at least 90% of the variance. The assumptions for these parametric tests seem to be verified, and are discussed with additional detail in Section 8.2.3. MANOVA appears to be more sensitive to implementation differences than the t-test, generally presenting more significant p-values. As would be expected, significant p-values from both MANOVA and t-tests are associated with clearer distinctions between samples in the PC scatter plots.

The concatenated output, $\widetilde{A}$, has less discriminatory power than individual outputs. This is especially clear for case II, where both the t-test and scatter plot do not suggest significant discrepancies. However, the MANOVA test is able to pick up the implementation difference, yielding a significant p-value. While in this case MANOVA on $\widetilde{A}$ answers to the question of whether the implementations are statistically equivalent, only a comprehensive analysis of individual outputs allows to diagnose how the model is affected by the implementation differences. For example, the number of PCs and the scatter plot indicate that the $\overline{E}^s$ output is by far the most affected in case II, something which is not possible to deduce by just analyzing $\widetilde{A}$. This consideration would also be difficult to infer from the empirical comparison data in Table 3, as the p-values are not much different from those associated with other outputs. As described in Section 5.1, outputs were range scaled prior to concatenation. We have experimented with other types of centering and scaling [58], and did not find major differences on how the proposed model-independent comparison method evaluates the concatenated output.

Results for the remaining size/set combinations, provided in Supplementary Tables S2.1–S2.8, are in accordance with the discussion thus far. The number of PCs required to explain 90% of the variance for cases II and III are generally smaller than for case I. Furthermore, the number of PCs for parameter set 2 is consistently lower in case II than in case III, attesting to what was observed in the model-dependent analysis: configurations compared in case II are more dissimilar than those in case III. This is also corroborated by the p-values, and even more clearly, by the PC scatter plots. For the majority of individual comparisons in cases II and III, MANOVA seems more sensitive to implementation differences than the t-test. However, there are some instances where the latter is able to perform better distinctions, as for example in case II of the 100@1 size/set combination. Concerning the concatenated output, it was generally adequate for detecting implementation differences, with MANOVA and/or the t-test yielding significant p-values. Nonetheless, this approach failed for case III of 100@2. Analyzing the overall results for all size/set combinations, it is also possible to conclude that, as model size increases, implementation differences become more pronounced for both parameter sets.

8.2.1 Assessing the t-test applied to individual PCs and the distribution of explained variance

Table 5 shows, for the first four PCs of the 400@1 combination, the percentage of variance explained by individual PCs, as well as the respective t-test p-values, before and after weighted Bonferroni correction. For case I, p-values for the first PC are never significant, but a few for the second PC are significant at the $\alpha = 0.05$ level. However, no p-values remain so after adjustment with the weighted Bonferroni correction. For case II, p-values for the first PC are usually the most significant, although this does not always hold, e.g., output $\widetilde{A}$. In this instance, only the 3^{rd}, 6^{th} and 15^{th} PC p-values are significant (the 3^{rd} and 6^{th} remain so at $\alpha = 0.05$ after correction), which is the reason for MANOVA catching the implementation difference (as shown in Table 4). For case III, the first PC p-value is always highly significant (before and after weighted Bonferroni correction), with the PC2–PC4 p-values showing no significance at all. The implementation differences in the first PC, captured by the t-test, seem sufficient for MANOVA, which considers all PCs simultaneously, to also spot the dissimilarities. The distribution of explained variance along PCs also reflects what was discussed in Section 5. More specifically, it is possible to observe that the first PC or PCs explain more variance in misaligned cases II and III than in case I (aligned).

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	% var. (PC1)	19.1	13.0	15.9	11.7	8.2	15.8	10.9
	% var. (PC2)	8.6	7.5	8.7	6.4	6.2	8.6	6.4
	% var. (PC3)	7.1	6.8	7.5	5.6	4.5	7.4	6.1
	% var. (PC4)	6.4	6.5	6.6	5.3	4.3	6.6	5.8
	t-test (PC1)	0.530	0.836	0.804	0.784	0.378	0.805	0.879
	t-test (PC2)	0.021	0.057	0.012	0.115	0.915	0.012	0.041
	t-test (PC3)	0.319	0.608	0.491	0.081	0.372	0.496	0.389
	t-test (PC4)	0.763	0.200	0.872	0.473	0.327	0.868	0.524
	t-test* (PC1)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC2)	0.240	0.757	0.135	1.000	1.000	0.136	0.633
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
II	% var. (PC1)	20.1	19.3	16.4	77.2	8.0	16.2	11.4
	% var. (PC2)	10.7	10.5	11.1	2.6	6.1	11.0	8.8
	% var. (PC3)	7.0	8.0	7.3	1.9	4.7	7.2	6.2
	% var. (PC4)	5.9	5.3	6.3	1.3	4.1	6.3	5.5
	t-test (PC1)	0.042	<1e-08	0.108	<1e-08	0.017	0.109	0.390
	t-test (PC2)	0.128	0.029	0.106	0.720	0.175	0.107	0.086
	t-test (PC3)	0.004	0.110	0.006	0.885	0.909	0.006	0.002
	t-test (PC4)	0.767	0.401	0.966	0.978	0.060	0.954	0.631
	t-test* (PC1)	0.209	<1e-08	0.661	<1e-08	0.211	0.669	1.000
	t-test* (PC2)	1.000	0.276	0.955	1.000	1.000	0.971	0.975
	t-test* (PC3)	0.055	1.000	0.079	1.000	1.000	0.077	0.028
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
III	% var. (PC1)	38.6	79.2	74.3	73.6	32.4	94.0	29.1
	% var. (PC2)	9.4	2.6	3.4	2.8	4.4	0.8	7.9
	% var. (PC3)	7.6	2.4	3.2	2.4	4.2	0.7	7.7
	% var. (PC4)	6.9	2.0	2.7	2.0	3.7	0.6	5.7
	t-test (PC1)	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	t-test (PC2)	0.183	0.966	0.792	0.662	0.979	0.878	0.792
	t-test (PC3)	0.734	0.906	0.826	0.976	0.777	0.951	0.754
	t-test (PC4)	0.602	0.800	0.587	0.827	0.979	0.875	0.555
	t-test* (PC1)	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000

Table 5 – Percentage of explained variance (% var.) and t-test p-values before (t-test) and after (t-test*) weighted Bonferroni correction of the first four PCs for model size 400, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

These observations generally remain valid when considering all tested size/set combinations (Tables S3.1–S3.8, provided as Supplementary material). For case I, in which the compared configurations are assumed to be aligned, a few significant PC1 p-values (at the $\alpha = 0.05$ level) stand-out for 100@1. Nonetheless, after the weighted Bonferroni correction, all t-test p-values from PC1–PC4 are non-significant for all size/set combinations. For case II, a global size/set analysis confirms that p-values for the first PC are usually the most significant for all outputs. The most notable exception occurs for the 200@1 combination, in which only the $\overline{E}^s$ output presents a significant p-value (at $\alpha = 0.01$, before and after weighted Bonferroni correction). However, this p-value is enough for the respective MANOVA (Table S2.3) to catch the difference directly in $\overline{E}^s$, and indirectly in the concatenated output. With very few exceptions, the PC2–PC4 p-values are not significant, especially after the weighted Bonferroni correction. Results for case III broadly confirm what was observed for 400@1, i.e., that the first PC p-value is always highly significant (before and after weighted Bonferroni correction), with the PC2–PC4 p-values showing little to no significance. The exception is 100@2, for which only the PC1 p-value for P^s is significant at $\alpha = 0.05$, but losing significance after the weighted Bonferroni correction. Nonetheless, dissimilarities between configurations 1 and 4 are detected by the MANOVA test (Table S2.2) in the majority of outputs due to misalignments in PCs other than the first (e.g., for output $\overline{C}$ these are visible for PC2 and PC3). Considering these results, it is possible to conclude that the p-value of the first PC is undoubtedly the most important for evaluating model alignment. Nonetheless, there are also instances, namely for smaller model sizes, where it is necessary to look at the MANOVA p-value in order to draw solid conclusions. The distribution of explained variance along PCs generally reflects the alignment of configurations, with less aligned ones having more variance explained by the first PC or PCs. Again the exception is 100@2, where the difference in explained variance does not change much from case to case.

8.2.2 Changing the variance for the selection of the number of PCs for MANOVA

The MANOVA p-values shown in Table 4 were obtained by requiring that the respective number of PCs explain at least 90% of the variance in the data. As this percentage must be predefined by the model developer, it is important to understand how the MANOVA results are affected when this value is changed. Table 6 shows, for the 400@1 instance, the number of PCs, as well as the associated MANOVA p-values, required to explain 30%, 50%, 70% and 90% of the variance. Three aspects can be highlighted from this table: 1) conclusions concerning configuration alignment do not change with different percentages of variance; 2) more variance seems to make the MANOVA test more sensitive in case II, although this trend is not very well-defined; 3) lower prespecified percentages of variance imply less PCs to explain it; in the limit, this makes the MANOVA test inapplicable if just one PC is required, as observed for several instances of cases II and III. This latter aspect is not crucial, because the t-test is the equivalent of a "univariate MANOVA" for two groups. Nonetheless, if the MANOVA p-values are deemed important for a specific comparison, it seems preferable to specify a higher percentage of variance to explain.

Results for the remaining size/set combinations, provided in Supplementary Tables S4.1–S4.8, generally follow the tendencies verified for 400@1, perhaps further blurring the slight trend observed in the relation between percentage of variance and MANOVA sensitivity.

8.2.3 Assumptions for the t and MANOVA parametric tests

As described in Section 5, the t and MANOVA tests make several assumptions about the underlying data. To confirm that the presented results are statistically valid, it is important to perform an overall survey of these assumptions, namely: 1) whether the data is normally distributed within groups; and, 2) if samples are drawn from populations with equal variances. The former assumption can be verified with the Shapiro–Wilk test for univariate normality or the Royston test for multivariate normality. The latter can be evaluated with the Bartlett test or Box's M test for the univariate and multivariate cases, respectively. Table 7 presents aggregated results for these tests for all size/set combinations.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs (30% var.)	3	4	3	5	6	3	5
	#PCs (50% var.)	6	8	6	10	13	7	9
	#PCs (70% var.)	11	15	12	19	25	12	19
	#PCs (90% var.)	24	32	25	36	43	26	39
	MNV (30% var.)	0.080	0.235	0.078	0.185	0.696	0.078	0.297
	MNV (50% var.)	0.267	0.514	0.278	0.020	0.530	0.311	0.614
	MNV (70% var.)	0.241	0.611	0.344	0.044	0.682	0.344	0.527
	MNV (90% var.)	0.528	0.258	0.548	0.105	0.746	0.577	0.704
II	#PCs (30% var.)	2	3	3	1	6	3	4
	#PCs (50% var.)	6	6	6	1	13	6	9
	#PCs (70% var.)	11	12	11	1	24	12	18
	#PCs (90% var.)	24	29	24	12	43	25	38
	MNV (30% var.)	0.038	≤1e-08	0.004	*	0.035	0.004	0.007
	MNV (50% var.)	2e-06	≤1e-08	3e-04	*	0.034	3e-04	3e-04
	MNV (70% var.)	4e-05	≤1e-08	0.002	*	0.049	0.001	2e-06
	MNV (90% var.)	≤1e-08	≤1e-08	5e-08	≤1e-08	0.467	1e-07	≤1e-08
III	#PCs (30% var.)	1	1	1	1	1	1	2
	#PCs (50% var.)	3	1	1	1	6	1	4
	#PCs (70% var.)	6	1	1	1	17	1	11
	#PCs (90% var.)	17	7	8	13	38	1	31
	MNV (30% var.)	*	*	*	*	*	*	≤1e-08
	MNV (50% var.)	≤1e-08	*	*	*	≤1e-08	*	≤1e-08
	MNV (70% var.)	≤1e-08	*	*	*	≤1e-08	*	≤1e-08
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	*	≤1e-08

Table 6 – Number of principal components (#PCs) and MANOVA test p-values (MNV) for several percentages of explained variance for model size 400, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replaced with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

More specifically, this table shows the percentage of p-values from the specified tests which fall within the non-significant ($p > 0.05$) and significant ($0.01 < p < 0.05$ and $p < 0.01$) intervals.

Test	$p > 0.05$	$0.01 < p < 0.05$	$p < 0.01$
SW (PC1)	94.77% (95.24%)	4.19% (4.46%)	1.04% (0.30%)
Bartlett (PC1)	92.65% (89.29%)	5.15% (8.93%)	2.20% (1.79%)
Royston	88.51%	6.52%	4.97%
Box's M	27.33%	16.15%	56.52%

Table 7 – Percentage of p-values yielded by the specified tests of assumptions which fall within the non-significant ($p > 0.05$) and significant ($0.01 < p < 0.05$ and $p < 0.01$) intervals. For the Shapiro–Wilk (SW) and Bartlett tests, percentages outside of parenthesis refer to tests of individual groups for each PC, while percentages within parentheses refer to tests of individual groups for the first PC. These percentages consider all tested size/set combinations, for $n = 30$ runs per configuration, and variance to explain prespecified at 90%.

The Shapiro–Wilk test does not reject univariate normality in approximately 95% of samples used for the t-tests. Of the remaining 5%, only 1% show significant rejection at the $\alpha = 0.01$ level. This percentage is even lower for PC1 samples, the most critical for evaluating configuration alignment using the univariate approach. Samples also seem to generally have similar variance, as denoted by the 90% of non-significant results for the Bartlett test. Multivariate normality is also verified by the Royston test in most instances. Since p-values are the probability of obtaining a result at least as extreme than what was actually observed (assuming that the null hypothesis is true), type I error rates are close to what would be expected, namely 95% for $p > 0.05$, 4% for $0.01 < p < 0.05$ and 1% for $p < 0.01$. As such, the assumptions tested by the Shapiro–Wilk, Bartlett and Royston tests seem to be generally verified. The same does not appear to hold for the MANOVA assumption of homogeneity of variance–covariance matrices, evaluated with Box's M test. However, as stated in Section 3, this test is extremely sensitive and MANOVA is resilient to violations in this assumption when samples have equal sizes, which is the case here.

8.3 Comparison performance with different sample sizes

Larger sample sizes make hypothesis test more powerful, i.e., more likely to reject the null hypothesis when it is false, thus decreasing the probability of committing a type II error. Here we are essentially interested in determining, within the discussed frameworks, if and how conclusions concerning configuration alignment change with sample size. For this purpose, the empirical and the proposed model-independent comparison methods were tested with samples of $n = 10$ and $n = 100$ replicates for the 400@1 combination. The comparison with samples of size 10 used the first 10 observations from the $n = 30$ runs per configuration setup, while the comparison with samples of size 100 used a new set of $n = 100$ runs per configuration. Results from this analysis are provided in Supplementary Tables S5.1 (empirical approach, 10 runs), S5.2 (model-independent approach, 10 runs), S6.1 (empirical approach, 100 runs) and S6.2 (model-independent approach, 100 runs).

Using samples of 10 observations, the empirical selection of FMs still manages to find configuration differences in cases II and III, although not as clearly as for $n = 30$. In case II, the difference is found on the steady-state mean of the P^w and $\overline{E}^s$ outputs. With $n = 30$, differences are detected in several more FMs. The proposed model-independent comparison approach is also able to spot differences in cases II and III, likewise detecting differences in the P^w and $\overline{E}^s$ outputs for case II.

For $n = 100$, both methods confirm what was observed for smaller sample sizes, although results are more pronounced, with very significant p-values for cases II and III. Two details emerge, however. First, for case II, the empirical approach hardly recognizes configuration differences for the mean predator energy output, $\overline{E}^w$. Such differences are exposed by the proposed model-independent method, both by

the t-test and MANOVA. Second, the model-independent method perceives a difference where there should not be any, namely in output $\overline{E}^s$ of case I, via the MANOVA test. The PC1 t-test p-value is not significant, as is the case for the vast majority of remaining PCs. For the 72 PCs (dimensions) considered by MANOVA, the t-test p-value is significant at $\alpha = 0.01$ in three cases: PC11, PC22 and PC48. After weighted Bonferroni correction, only the PC11 p-value, associated with 2.4% of the variance, remains significant, with a value of 0.006. Since these PCs, and PC11 in particular, account for a small amount of output variability, conclusions concerning output similarity (and consequently, model alignment), do not necessarily change. In practice this result is not totally unexpected, since the increase in power (due to a larger sample size) together with a MANOVA test on many PCs or dimensions, leads to a higher probability of detecting smaller or irrelevant differences. Overall, what seems clear for a sample of $n = 100$ replicates, is that a t-test on the first PC of each individual output (or even the concatenated output) is enough to pick up implementation differences.

9 Discussion

The empirical approach and the proposed model-independent method were able to identify the changes introduced in configurations 3 and 4, even though these were of different nature. The change introduced in configuration 3, case II, confirms the sensitivity of this type of models: a trivial implementation discrepancy led to statistically different behavior. In turn, the change in configuration 4, case III, was expected to be noticed, since it implied the alteration of a model parameter. Nonetheless, in specific conditions, namely for parameter set 2, configuration 3 presented more dissimilar behavior than configuration 4, when compared to configuration 1. Since parameter set 2 generates more agents during simulations (as described in Section 6.2), the effect of not shuffling the agent list in configuration 3, case II, may be more pronounced. Additionally, the c_r parameter is larger in parameter set 2 than in parameter set 1. As such, the smaller relative difference of decreasing c_r by one unit in configuration 4 may lead to less significant discrepancies in case III.

In the context of a smaller sample size, namely for $n = 10$ runs per configuration, both methods were capable of discerning differences in cases II and III. Differentiation was better with increasing number of replicates, i.e., when going from $n = 10$, then $n = 30$, and finally $n = 100$. In practice, however, more replicates did not yield different conclusions regarding output and configuration alignment.

Globally, results for the proposed model-independent method were in accordance with the empirical approach on an per output basis. The method was also able to expose a few implementation differences not detected when empirically selecting FMs, and, in other instances, it identified these differences more conclusively. As such, the proposed method was shown to be a more direct, and often more effective alternative to the empirical and model-dependent selection of FMs.

The p-values yielded by the t and MANOVA tests offer an objective assessment concerning the alignment of individual outputs. While in most instances the two tests generated the same decision, there were situations in which results differed. In case of alignment, the t-test on the first PC never produced p-values significant at $\alpha = 0.01$, even prior to weighted Bonferroni correction. MANOVA, however, signaled a few false positives for the $\overline{E}^s$ output, especially in the $n = 100$ observations per sample case, which also affected the $\widetilde{A}$ output. A careful analysis revealed the cause to be a misalignment in one or more PCs, which in themselves did not account for much output variance, but were sufficient to provoke a type I error in the MANOVA test. In case of configuration misalignment, MANOVA was generally more sensitive than the PC1 t-test, and with the exception of 100@1, flagged a few misalignments unnoticed by the t-test.

The concatenated output provides a single answer on whether two or more model implementations are aligned or not. In terms of false positives on $\widetilde{A}$, i.e., reporting a misalignment where none is expected, this was observed once by a MANOVA test for the $n = 100$ runs per configuration instance. The t-test, however, did not flag any configuration difference. False negatives, i.e., failure to detect output mismatch when it is known to exist, occurred a few times for both MANOVA and the t-test, and were essentially associated with smaller model sizes. Additionally, neither test was able to detect

differences on $\widetilde{A}$ in case II for the $n = 10$ runs per configuration instance.

10 Recommendations on using the model-independent method

Given the results presented and discussed in the previous sections, a number of recommendations on using the proposed model-independent method can be made.

The t and MANOVA tests should both be performed, and in case of disagreement, it is important to find the cause. More specifically, 1) if the PC1 t-test p-value is significant (and the MANOVA p-value is not), it should be adjusted with the weighted Bonferroni correction, and considered instead; 2) if the MANOVA p-value is significant (and the PC1 t-test p-value is not), the t-test should be performed on other PCs to find which ones are likely to be influencing the MANOVA result. A decision on alignment should consider how much variance the culprit PCs explain. As an alternative to performing the t-test on PC1 and MANOVA on the number of PCs which explain a minimum percentage of variance, it is also possible to perform the t-test on all PCs individually, and consider the resulting p-values after weighted Bonferroni correction. This also has the advantage of automatically privileging the most important PCs, i.e., the ones that explain more variance. Furthermore, if samples do not appear to be drawn from a normal distribution, a non-parametric test can be used instead of the t-test (e.g., Mann–Whitney or Kolmogorov–Smirnov).

The prespecified variance determines the number of PCs used in the MANOVA test. While conclusions concerning configuration alignment do not appear to change dramatically with different percentages of variance, it seems preferable to specify more variance than less, otherwise one may not be able to use MANOVA. A minimum variance of 90% seems appropriate for this purpose, but this may vary with models, outputs characteristics and parameterizations. Another possibility is to use a fixed number of PCs for MANOVA instead of specifying a percentage of variance to explain. A potential problem with this approach is that MANOVA could eventually consider very different percentages of variance from output to output and/or case to case.

The eigenvalue structure, reflected in the number of PCs required to explain a prespecified percentage of variance, as well as in the variance explained by the first PC or PCs, is also an interesting comparison metric to consider. For aligned outputs, in contrast with misaligned ones, more PCs are required to explain the same amount of variance, and less variance is explained by the first PC or PCs. However, the information provided by this metric in general, and by the number of PCs in particular, is only useful when assessing multiple cases, for which it can be relatively compared.

The scatter plot of the first two PC dimensions is also useful for quickly assessing possible output misalignments. If the outputs have a clear mismatch, points associated with runs from different implementations form visibly distinct groups. Although this technique is basically a type of face validation, it provides a useful intuition on the alignment of two or more implementations for any given output.

Finally, output concatenation also provides a practical and rapid way of assessing model alignment, but should not be fully relied upon to make a final decision. Additionally, the analysis of individual outputs allows to diagnose which model aspects are most affected by possible mismatches, something which the concatenated output does not permit.

Altogether, no single metric of the proposed method provides a definitive answer, and decisions concerning model alignment should be taken based on the different types of information it produces.

11 Conclusions

A technique for comparing simulation models in a model-independent fashion was presented in this paper. The method dispenses the need for selecting model-specific FMs, and is appropriate for comparing multiple implementations of a conceptual model or replications of an existing model. The technique was validated against the commonly used approach of empirically selecting output statistical summaries using the PPHPC ABM on a number of model sizes and parameter sets. Results showed that

the proposed method offered similar conclusions concerning output and model alignment, providing stronger misalignment detection power in some situations.

Although the proposed framework is presented in the context of simulation output comparison, it is a generic statistical comparison methodology. As such, it can be employed in circumstances where observations yield a large number of correlated variables [60].

Acknowledgments

This work was supported by the Fundação para a Ciência e a Tecnologia (FCT) projects UID/EEA/50009/2013 and UID/MAT/04561/2013, and partially funded with grant SFRH/BD/48310/2008, also from FCT.

References

[1] N. Boccara, Modeling Complex Systems, Graduate Texts in Contemporary Physics series, Springer-Verlag New York, 2004. doi:10.1007/b97378.

[2] Y. Bar-Yam, Dynamics of complex systems, Vol. 213, Addison-Wesley Reading, MA, 1997.

[3] A. Bagdasaryan, Discrete dynamic simulation models and technique for complex control systems, Simulation Modelling Practice and Theory 19 (4) (2011) 1061–1087. doi:10.1016/j.simpat.2010.12.010.

[4] V. Grimm, U. Berger, D. DeAngelis, J. Polhill, J. Giske, S. Railsback, The ODD protocol: A review and first update, Ecological Modelling 221 (23) (2010) 2760–2768. doi:10.1016/j.ecolmodel.2010.08.019.

[5] A. S. Shirazi, T. Davison, S. von Mammen, J. Denzinger, C. Jacob, Adaptive agent abstractions to speed up spatial agent-based simulations, Simulation Modelling Practice and Theory 40 (2014) 144–160. doi:10.1016/j.simpat.2013.09.001.

[6] J.-S. Lee, T. Filatova, A. Ligmann-Zielinska, B. Hassani-Mahmooei, F. Stonedahl, I. Lorscheid, A. Voinov, J. G. Polhill, Z. Sun, D. C. Parker, The complexities of agent-based modeling output analysis, Journal of Artificial Societies and Social Simulation 18 (4) (2015) 4. doi:10.18564/jasss.2897.
URL http://jasss.soc.surrey.ac.uk/18/4/4.html

[7] U. Merlone, M. Sonnessa, P. Terna, Horizontal and vertical multiple implementations in a model of industrial districts, Journal of Artificial Societies and Social Simulation 11 (2) (2008) 5.
URL http://jasss.soc.surrey.ac.uk/11/2/5.html

[8] O. Will, R. Hegselmann, A replication that failed – on the computational model in 'Michael W. Macy and Yoshimichi Sato: Trust, cooperation and market formation in the U.S. and Japan. Proceedings of the National Academy of Sciences, May 2002', Journal of Artificial Societies and Social Simulation 11 (3) (2008) 3.
URL http://jasss.soc.surrey.ac.uk/11/3/3.html

[9] U. Wilensky, W. Rand, Making models match: replicating an agent-based model, Journal of Artificial Societies and Social Simulation 10 (4) (2007) 2.
URL http://jasss.soc.surrey.ac.uk/10/4/2.html

[10] S. Robinson, R. E. Nance, R. J. Paul, M. Pidd, S. J. Taylor, Simulation model reuse: definitions, benefits and obstacles, Simulation Modelling Practice and Theory 12 (7–8) (2004) 479–494. doi:10.1016/j.simpat.2003.11.006.

[11] J. C. Thiele, V. Grimm, Replicating and breaking models: good for you and good for ecology, Oikos 124 (6) (2015) 691–696. doi:10.1111/oik.02170.

[12] B. Edmonds, D. Hales, Replication, replication and replication: Some hard lessons from model alignment, Journal of Artificial Societies and Social Simulation 6 (4) (2003) 11.
URL http://jasss.soc.surrey.ac.uk/6/4/11.html

[13] N. David, Validating simulations, in: Simulating Social Complexity, Understanding Complex Systems, Springer Berlin Heidelberg, 2013, pp. 135–171. doi:10.1007/978-3-540-93813-2_8.

[14] N. Fachada, V. V. Lopes, R. C. Martins, A. C. Rosa, Parallelization strategies for spatial agent-based models, International Journal of Parallel Programming (2016) 1–33doi:10.1007/s10766-015-0399-9.

[15] H. R. Parry, M. Bithell, Large scale agent-based modelling: A review and guidelines for model scaling, in: A. J. Heppenstall, A. T. Crooks, L. M. See, M. Batty (Eds.), Agent-Based Models of Geographical Systems, Springer Netherlands, 2012, pp. 271–308. doi:10.1007/978-90-481-8927-4_14.

[16] R. D. Peng, Reproducible research in computational science, Science 334 (6060) (2011) 1226–1227. doi:10.1126/science.1213847.

[17] R. Axtell, R. Axelrod, J. M. Epstein, M. D. Cohen, Aligning simulation models: a case study and results, Computational and Mathematical Organization Theory 1 (2) (1996) 123–141. doi:10.1007/BF01299065.

[18] B. Müller, S. Balbi, C. M. Buchmann, L. de Sousa, G. Dressler, J. Groeneveld, C. J. Klassert, Q. B. Le, J. D. Millington, H. Nolzen, D. C. Parker, J. G. Polhill, M. Schlüter, J. Schulze, N. Schwarz, Z. Sun, P. Taillandier, H. Weise, Standardised and transparent model descriptions for agent-based models: Current status and prospects, Environmental Modelling & Software 55 (2014) 156–163. doi:10.1016/j.envsoft.2014.01.029.

[19] I. Jolliffe, Principal component analysis, 2nd Edition, Springer Series in Statistics, Springer, 2002. doi:10.1007/b98835.

[20] N. Fachada, V. V. Lopes, R. C. Martins, A. C. Rosa, Towards a standard model for research in agent-based modeling and simulation, PeerJ Computer Science 1 (2015) e36. doi:10.7717/peerj-cs.36.

[21] K. M. Carley, Computational organizational science and organizational engineering, Simulation Modelling Practice and Theory 10 (5) (2002) 253—-269. doi:10.1016/S1569-190X(02)00119-3.

[22] N. Fachada, V. V. Lopes, A. Rosa, Simulating antigenic drift and shift in influenza A, in: Proc. of the 2009 ACM symposium on Applied Computing, SAC '09, ACM, New York, NY, USA, 2009, pp. 2093–2100. doi:10.1145/1529282.1529744.
URL http://doi.acm.org/10.1145/1529282.1529744

[23] O. Balci, R. G. Sargent, Validation of simulation models via simultaneous confidence intervals, American Journal of Mathematical and Management Sciences 4 (3–4) (1984) 375–406. doi:10.1080/01966324.1984.10737151.

[24] D. Miodownik, B. Cartrite, R. Bhavnani, Between replication and docking: "Adaptive agents, political institutions, and civic traditions" revisited, Journal of Artificial Societies and Social Simulation 13 (3) (2010) 1.
URL http://jasss.soc.surrey.ac.uk/13/3/1.html

[25] W. Radax, B. Rengs, Prospects and pitfalls of statistical testing: Insights from replicating the demographic prisoner's dilemma, Journal of Artificial Societies and Social Simulation 13 (4) (2010) 1.
URL http://jasss.soc.surrey.ac.uk/13/4/1.html

[26] R. G. Sargent, A new statistical procedure for validation of simulation and stochastic models, Tech. Rep. SYR-EECS-2010-06, Department of Electrical Engineering and Computer Science, Syracuse University (Nov. 2010).

[27] A. M. Law, Simulation Modeling and Analysis, 5th Edition, McGraw-Hill, Columbus, OH, USA, 2015.

[28] F. J. Massey Jr., The Kolmogorov-Smirnov test for goodness of fit, Journal of the American Statistical Association 46 (253) (1951) 68–78. doi:10.1080/01621459.1951.10500769.

[29] J. D. Gibbons, S. Chakraborti, Nonparametric statistical inference, 5th Edition, Statistics: Textbooks & Monographs, Chapman and Hall/CRC, Boca Raton, FL, USA, 2010.

[30] R. Bhavnani, Adaptive agents, political institutions and civic traditions in modern Italy, Journal of Artificial Societies and Social Simulation 6 (4) (2003) 1.
URL http://jasss.soc.surrey.ac.uk/6/4/1.html

[31] V. Dergachev, M. Garces, Ps-i, SourceForge (2013).
URL http://ps-i.sourceforge.net/

[32] J. M. Epstein, Zones of cooperation in demographic prisoner's dilemma, Complexity 4 (2) (1998) 36–48. doi:10.1002/(SICI)1099-0526(199811/12)4:2<36::AID-CPLX9>3.0.CO;2-Z.

[33] S. Alberts, M. K. Keenan, R. M. D'Souza, G. An, Data-parallel techniques for simulating a mega-scale agent-based model of systemic inflammatory response syndrome on graphics processing units, Simulation 88 (8) (2011) 895–907. doi:10.1177/0037549711425180.

[34] Nvidia, CUDA Toolkit Documentation v7.0, Nvidia Corporation (Mar. 2015).

[35] G. An, Toy infection model, NetLogo User Community Models (Jul. 2007).
URL http://ccl.northwestern.edu/netlogo/models/community/Toy_Infection_Model

[36] U. Wilensky, NetLogo (1999).
URL http://ccl.northwestern.edu/netlogo/

[37] W. H. Kruskal, W. A. Wallis, Use of ranks in one-criterion variance analysis, Journal of the American Statistical Association 47 (260) (1952) 583–621. doi:10.1080/01621459.1952.10483441.

[38] D. C. Montgomery, G. C. Runger, Applied statistics and probability for engineers, sixth Edition, John Wiley & Sons, 2014.

[39] W. J. Krzanowski, Principles of Multivariate Analysis: A User's Perspective, Oxford University Press, New York, USA, 1988.

[40] B. G. Tabachnick, L. S. Fidell, Using multivariate statistics, sixth Edition, Pearson, 2013.

[41] M. J. Anderson, A new method for non-parametric multivariate analysis of variance, Austral Ecology 26 (1) (2001) 32–46. doi:10.1111/j.1442-9993.2001.01070.pp.x.

[42] K. R. Clarke, Non-parametric multivariate analyses of changes in community structure, Australian Journal of Ecology 18 (1) (1993) 117–143. doi:10.1111/j.1442-9993.1993.tb00438.x.
URL http://dx.doi.org/10.1111/j.1442-9993.1993.tb00438.x

[43] P. W. Mielke Jr, K. J. Berry, E. S. Johnson, Multi-response permutation procedures for a priori classifications, Communications in Statistics - Theory and Methods 5 (14) (1976) 1409–1424. doi:10.1080/03610927608827451.

[44] P. R. Rosenbaum, An exact distribution-free test comparing two multivariate distributions based on adjacency, Journal of the Royal Statistical Society: Series B (Statistical Methodology) 67 (4) (2005) 515–530. doi:10.1111/j.1467-9868.2005.00513.x.

[45] L. Baringhaus, C. Franz, On a new multivariate two-sample test, Journal of Multivariate Analysis 88 (1) (2004) 190–206. doi:10.1016/s0047-259x(03)00079-4.

[46] C. A. Boneau, The effects of violations of assumptions underlying the t test, Psychological Bulletin 57 (1) (1960) 49–64. doi:10.1037/h0041412.

[47] O. Balci, R. G. Sargent, Validation of multivariate response models using Hotelling's two-sample T^2 test, Simulation 39 (6) (1982) 185–192. doi:10.1177/003754978203900602.

[48] S. S. Shapiro, M. B. Wilk, An analysis of variance test for normality (complete samples), Biometrika 52 (3/4) (1965) 591–611. doi:10.2307/2333709.

[49] J. Royston, An extension of Shapiro and Wilk's W test for normality to large samples, Journal of the Royal Statistical Society. Series C, Applied Statistics 31 (2) (1982) 115–124. doi:10.2307/2347973.

[50] M. S. Bartlett, Properties of sufficiency and statistical tests, Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences 160 (901) (1937) 268–282. doi:10.1098/rspa.1937.0109.

[51] G. E. P. Box, A general distribution theory for a class of likelihood criteria, Biometrika 36 (3/4) (1949) 317–346. doi:10.2307/2332671.

[52] J. P. Shaffer, Multiple hypothesis testing, Annual Review of Psychology 46 (1995) 561–584. doi:10.1146/annurev.ps.46.020195.003021.

[53] T. V. Perneger, What's wrong with bonferroni adjustments, BMJ 316 (7139) (1998) 1236–1238. doi:10.1136/bmj.316.7139.1236.

[54] M. Matsumoto, T. Nishimura, Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator, ACM Transactions on Modeling and Computer Simulation 8 (1) (1998) 3–30. doi:10.1145/272991.272995.

[55] S. M. Sanchez, ABC's of output analysis, in: Simulation Conference Proceedings, 1999 Winter, Vol. 1, IEEE, Phoenix, AZ, USA, 1999, pp. 24–32. doi:10.1109/WSC.1999.823048.

[56] N. Fachada, M. A. T. Figueiredo, V. V. Lopes, R. C. Martins, A. C. Rosa, Spectrometric differentiation of yeast strains using minimum volume increase and minimum direction change clustering criteria, Pattern Recognition Letters 45 (2014) 55–61. doi:10.1016/j.patrec.2014.03.008.

[57] R. Rosenthal, D. B. Rubin, Ensemble-adjusted p values, Psychological Bulletin 94 (3) (1983) 540–541. doi:10.1037/0033-2909.94.3.540.

[58] R. A. Berg, H. C. Hoefsloot, J. A. Westerhuis, A. K. Smilde, M. J. Werf, Centering, scaling, and transformations: improving the biological information content of metabolomics data, BMC Genomics 7 (1) (2006) 142. doi:10.1186/1471-2164-7-142.

[59] N. Fachada, V. V. Lopes, R. C. Martins, A. C. Rosa, SimOutUtils - utilities for analyzing time series simulation output, Journal of Open Research Software 4 (1) (2016) e38. doi:10.5334/jors.110.

[60] N. Fachada, J. Rodrigues, V. V. Lopes, R. C. Martins, A. C. Rosa, micompr: An R package for multivariate independent comparison of observations, The R Journal[IN PRESS].

Supplementary tables

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.151	0.620	0.016	0.028	0.343	0.138
	arg max	0.424	0.853	0.033	0.106	0.264	0.005
	min	0.007	0.054	0.082	3e−5	0.152	0.014
	arg min	0.623	0.842	0.014	1.000	0.168	0.140
	$\overline{X}^{ss}$	0.446	0.292	0.212	0.838	0.945	0.211
	S^{ss}	0.143	0.226	0.131	0.109	0.237	0.130
II	max	0.924	0.017	4e−5	4e−6	0.250	0.974
	arg max	2e−4	0.355	0.223	0.002	0.147	4e−5
	min	8e−5	0.009	0.830	2e−5	0.025	6e−5
	arg min	0.043	0.544	0.001	1.000	0.795	0.591
	$\overline{X}^{ss}$	0.005	$<$ 1e−8	0.325	$<$ 1e−8	0.337	0.323
	S^{ss}	0.004	0.007	0.005	0.005	0.204	0.005
III	max	$<$ 1e−8	$<$ 1e−8	2e−5	0.029	0.303	$<$ 1e−8
	arg max	0.015	0.652	$<$ 1e−8	0.876	0.399	2e−4
	min	0.013	6e−5	0.463	0.001	0.094	$<$ 1e−8
	arg min	0.042	0.636	0.001	1.000	0.378	$<$ 1e−8
	$\overline{X}^{ss}$	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8
	S^{ss}	9e−8	$<$ 1e−8	5e−6	$<$ 1e−8	0.002	0.003

Table S1.1 P-values for the empirical selection of FMs for model size 100, parameter set 1, and $n = 30$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{ss}$ and S^{ss} statistics, and with the Mann-Whitney U test for arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.012	0.610	0.561	0.473	0.566	0.969
	arg max	0.792	0.085	0.273	0.602	0.292	0.050
	min	0.429	0.065	0.156	0.015	0.550	0.546
	arg min	0.636	0.894	0.209	0.145	1.000	0.424
	$\overline{X}^{\text{ss}}$	0.675	0.926	0.960	0.970	0.594	0.973
	S^{ss}	0.615	0.776	0.781	0.898	0.789	0.782
II	max	< 1e−8	2e−6	0.001	0.087	1e−6	9e−5
	arg max	0.014	0.047	0.554	0.024	0.350	0.542
	min	0.009	0.630	8e−8	< 1e−8	0.898	0.001
	arg min	0.112	0.542	0.096	9e−5	0.392	0.492
	$\overline{X}^{\text{ss}}$	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8
	S^{ss}	0.090	0.058	0.699	0.870	0.028	0.706
III	max	< 1e−8	0.009	0.040	0.175	0.155	< 1e−8
	arg max	0.893	0.745	0.824	0.701	0.059	0.402
	min	0.147	0.058	0.616	0.975	0.503	0.009
	arg min	0.069	0.917	0.152	0.274	0.459	0.982
	$\overline{X}^{\text{ss}}$	5e−5	< 1e−8	< 1e−8	0.762	< 1e−8	< 1e−8
	S^{ss}	0.010	0.011	0.026	0.010	0.037	0.277

Table S1.2 P-values for the empirical selection of FMs for model size 100, parameter set 2, and $n = 30$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{\text{ss}}$ and S^{ss} statistics, and with the Mann-Whitney U test for arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.336	0.372	0.943	0.924	0.320	0.392
	arg max	0.524	0.267	1.000	0.625	1.000	0.917
	min	0.821	0.069	0.322	0.081	0.268	0.986
	arg min	0.391	0.641	0.894	1.000	0.111	0.654
	$\overline{X}^{\text{ss}}$	0.383	0.924	0.339	0.252	0.075	0.336
	S^{ss}	0.173	0.408	0.162	0.879	0.540	0.162
II	max	0.003	2e−7	0.349	6e−7	0.236	0.013
	arg max	0.003	0.446	0.764	0.005	0.464	0.073
	min	0.274	0.870	0.007	3e−4	0.907	0.258
	arg min	0.197	0.906	0.060	1.000	0.853	0.654
	$\overline{X}^{\text{ss}}$	0.617	< 1e−8	0.001	< 1e−8	0.788	0.001
	S^{ss}	0.090	0.078	0.124	0.029	0.646	0.123
III	max	< 1e−8	< 1e−8	< 1e−8	< 1e−8	5e−4	< 1e−8
	arg max	0.095	0.912	< 1e−8	0.640	0.982	0.005
	min	0.155	0.004	3e−6	0.021	0.908	< 1e−8
	arg min	0.034	0.111	0.005	1.000	0.864	< 1e−8
	$\overline{X}^{\text{ss}}$	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8
	S^{ss}	1e−8	< 1e−8	9e−7	6e−8	1e−5	0.001

Table S1.3 P-values for the empirical selection of FMs for model size 200, parameter set 1, and $n = 30$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{\text{ss}}$ and S^{ss} statistics, and with the Mann-Whitney U test for arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.499	0.992	0.578	3e−4	0.819	0.815
	arg max	0.068	0.354	0.911	0.387	0.738	0.472
	min	0.132	0.602	0.796	0.387	0.314	0.575
	arg min	0.157	0.777	0.869	0.643	0.461	0.749
	$\overline{X}^{\text{ss}}$	0.448	0.318	0.475	0.041	0.139	0.472
	S^{ss}	0.543	0.754	0.672	0.481	0.987	0.670
II	max	< 1e−8	< 1e−8	9e−6	0.001	< 1e−8	< 1e−8
	arg max	2e−8	0.784	1e−4	0.267	0.598	0.256
	min	0.036	0.855	< 1e−8	< 1e−8	0.886	1e−5
	arg min	0.061	0.497	0.887	< 1e−8	0.068	8e−5
	$\overline{X}^{\text{ss}}$	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8
	S^{ss}	0.093	0.135	0.797	0.219	0.099	0.804
III	max	< 1e−8	1e−8	3e−5	3e−6	4e−5	< 1e−8
	arg max	0.246	0.667	0.079	0.164	0.431	1e−6
	min	0.041	0.132	0.001	0.652	0.980	2e−6
	arg min	0.118	0.011	0.398	0.001	0.366	0.109
	$\overline{X}^{\text{ss}}$	< 1e−8	< 1e−8	< 1e−8	0.001	< 1e−8	< 1e−8
	S^{ss}	0.271	0.417	0.723	0.804	0.421	0.605

Table S1.4 P-values for the empirical selection of FMs for model size 200, parameter set 2, and $n = 30$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{\text{ss}}$ and S^{ss} statistics, and with the Mann-Whitney U test for arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.213	0.288	0.011	0.231	0.774	0.086
	arg max	0.733	0.802	0.056	0.858	0.284	0.279
	min	0.002	0.088	0.094	0.275	0.076	0.011
	arg min	0.048	0.350	0.279	1.000	0.091	1.000
	$\overline{X}^{ss}$	0.457	0.905	0.546	0.833	0.049	0.551
	S^{ss}	0.242	0.282	0.285	0.474	0.109	0.285
II	max	7e−6	1e−7	2e−4	2e−7	0.322	0.004
	arg max	1e−4	0.540	0.230	0.703	0.052	1e−4
	min	1e−5	4e−4	0.009	0.508	0.017	0.001
	arg min	0.712	0.906	0.001	1.000	0.009	1.000
	$\overline{X}^{ss}$	0.206	< 1e−8	6e−8	< 1e−8	0.938	6e−8
	S^{ss}	0.735	0.324	0.712	0.218	0.688	0.713
III	max	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8
	arg max	2e−5	0.115	< 1e−8	0.134	0.332	1e−5
	min	7e−8	< 1e−8	< 1e−8	0.582	0.003	< 1e−8
	arg min	2e−5	0.070	2e−5	1.000	0.004	< 1e−8
	$\overline{X}^{ss}$	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8
	S^{ss}	3e−7	7e−8	5e−5	4e−8	0.013	0.013

Table S1.5 P-values for the empirical selection of FMs for model size 400, parameter set 1, and $n = 30$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{ss}$ and S^{ss} statistics, and with the Mann-Whitney U test for arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.840	0.982	0.933	0.552	0.132	0.878
	arg max	0.742	0.509	0.887	0.689	0.126	0.098
	min	0.255	0.944	0.289	0.509	0.138	0.956
	arg min	0.161	0.473	0.081	0.132	0.004	0.470
	$\overline{X}^{\text{ss}}$	0.528	0.469	0.379	0.326	0.786	0.362
	S^{ss}	0.631	0.634	0.685	0.669	0.657	0.684
II	max	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	0.169	$<$ 1e−8	$<$ 1e−8
	arg max	$<$ 1e−8	0.423	$<$ 1e−8	0.649	0.161	0.207
	min	0.019	0.173	$<$ 1e−8	$<$ 1e−8	0.056	$<$ 1e−8
	arg min	0.006	0.091	0.988	$<$ 1e−8	0.055	$<$ 1e−8
	$\overline{X}^{\text{ss}}$	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8
	S^{ss}	0.135	0.289	0.871	0.245	0.258	0.870
III	max	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8
	arg max	0.035	0.315	2e−6	0.100	0.001	1e−7
	min	0.038	0.236	$<$ 1e−8	0.997	0.488	$<$ 1e−8
	arg min	0.022	0.442	0.020	0.001	0.190	1e−6
	$\overline{X}^{\text{ss}}$	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	2e−6	$<$ 1e−8	$<$ 1e−8
	S^{ss}	0.164	0.110	0.043	0.027	0.135	0.006

Table S1.6 P-values for the empirical selection of FMs for model size 400, parameter set 2, and $n = 30$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{\text{ss}}$ and S^{ss} statistics, and with the Mann-Whitney U test for arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.371	0.729	0.001	0.300	0.834	0.125
	arg max	1.000	0.813	1.000	0.406	0.407	0.618
	min	0.047	0.434	0.177	0.042	0.789	0.016
	arg min	0.001	0.448	0.463	1.000	0.945	1.000
	$\overline{X}^{ss}$	0.782	0.075	0.820	0.041	0.080	0.818
	S^{ss}	0.111	0.116	0.161	0.131	0.434	0.159
II	max	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	0.168	3e−8
	arg max	$<$ 1e−8	0.947	1.000	0.464	0.143	9e−5
	min	4e−6	6e−8	3e−8	0.045	0.348	$<$ 1e−8
	arg min	0.094	0.001	2e−6	1.000	0.916	1.000
	$\overline{X}^{ss}$	0.005	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	0.018	$<$ 1e−8
	S^{ss}	0.580	0.398	0.553	0.431	0.593	0.552
III	max	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8
	arg max	1e−8	$<$ 1e−8	$<$ 1e−8	0.647	0.145	$<$ 1e−8
	min	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	0.063	0.001	$<$ 1e−8
	arg min	1e−4	6e−7	$<$ 1e−8	1.000	0.456	$<$ 1e−8
	$\overline{X}^{ss}$	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8	$<$ 1e−8
	S^{ss}	1e−5	2e−6	0.001	4e−6	0.049	0.188

Table S1.7 P-values for the empirical selection of FMs for model size 800, parameter set 1, and $n = 30$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{ss}$ and S^{ss} statistics, and with the Mann-Whitney U test for arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.913	0.018	0.178	0.444	0.347	0.194
	arg max	1.000	0.203	0.774	1.000	0.068	0.510
	min	1.000	0.024	0.728	8e−6	0.053	0.182
	arg min	1.000	0.047	0.876	0.974	0.954	0.397
	$\overline{X}^{ss}$	0.638	0.515	0.028	0.136	0.261	0.029
	S^{ss}	0.718	0.609	0.639	0.696	0.646	0.642
II	max	< 1e−8	< 1e−8	< 1e−8	0.379	< 1e−8	< 1e−8
	arg max	< 1e−8	0.006	< 1e−8	0.002	0.065	0.249
	min	0.108	1e−5	< 1e−8	< 1e−8	0.001	< 1e−8
	arg min	0.042	0.007	0.006	< 1e−8	0.081	< 1e−8
	$\overline{X}^{ss}$	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8
	S^{ss}	0.619	0.773	0.469	0.022	0.977	0.476
III	max	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8
	arg max	2e−5	0.012	< 1e−8	0.001	1.000	< 1e−8
	min	1.000	0.911	< 1e−8	0.178	0.104	< 1e−8
	arg min	1.000	0.239	0.605	2e−4	0.132	< 1e−8
	$\overline{X}^{ss}$	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8	< 1e−8
	S^{ss}	0.775	0.783	0.337	0.317	0.924	0.084

Table S1.8 P-values for the empirical selection of FMs for model size 800, parameter set 2, and $n = 30$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{ss}$ and S^{ss} statistics, and with the Mann-Whitney U test for arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	26	32	26	37	43	27	38
	MNV	0.852	0.509	0.851	0.966	0.285	0.879	0.486
	t-test	0.036	0.057	0.049	0.023	0.503	0.049	0.077
	PCS	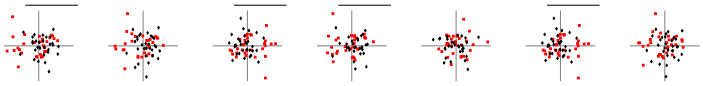						
II	#PCs	25	32	25	34	43	26	38
	MNV	0.023	4e-05	0.276	≤1e-08	0.448	0.256	0.496
	t-test	0.001	0.001	0.003	≤1e-08	0.956	0.003	0.003
	PCS	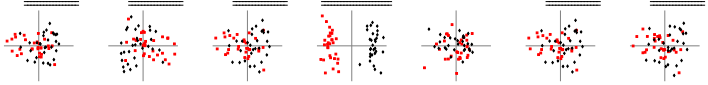						
III	#PCs	22	26	21	32	42	15	35
	MNV	≤1e-08	≤1e-08	≤1e-08	≤1e-08	0.001	≤1e-08	6e-06
	t-test	6e-08	≤1e-08	≤1e-08	≤1e-08	0.003	≤1e-08	≤1e-08
	PCS	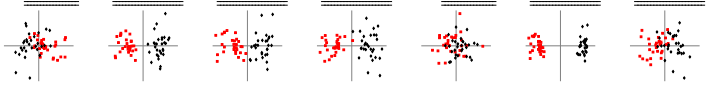						

Table S2.1 Model-independent comparison for model size 100, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) hypothesis tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replaced with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	20	19	15	29	31	15	22
	MNV	0.116	0.141	0.175	0.085	0.121	0.175	0.125
	t-test	0.108	0.128	0.106	0.111	0.067	0.106	0.099
	PCS	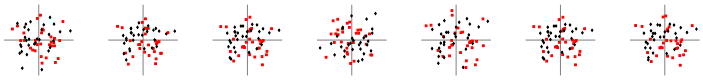						
II	#PCs	20	19	15	27	29	15	22
	MNV	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	5e-06
	t-test	<1e-08	1e-06	3e-08	<1e-08	<1e-08	3e-08	2e-07
	PCS	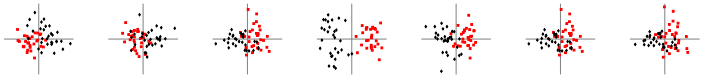						
III	#PCs	21	18	15	29	30	15	22
	MNV	0.006	0.001	0.110	0.139	2e-04	<1e-08	0.147
	t-test	0.037	0.472	0.322	0.324	0.054	0.116	0.346
	PCS							

Table S2.2 Model-independent comparison for model size 100, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) hypothesis tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replaced with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	25	32	25	36	43	26	38
	MNV	0.287	0.214	0.157	0.123	0.699	0.187	0.356
	t-test	0.783	0.810	0.895	0.973	0.830	0.895	0.979
	PCS							
II	#PCs	24	31	25	27	43	25	38
	MNV	0.022	≤1e-08	0.276	≤1e-08	0.778	0.270	0.038
	t-test	0.616	0.056	0.481	≤1e-08	0.773	0.480	0.534
	PCS							
III	#PCs	21	19	17	27	41	6	35
	MNV	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	3e-08
	t-test	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	PCS							

Table S2.3 Model-independent comparison for model size 200, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) hypothesis tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replaced with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	21	19	15	29	31	15	21
	MNV	0.138	0.264	0.322	0.342	0.652	0.322	0.244
	t-test	0.192	0.294	0.280	0.985	0.378	0.280	0.240
	PCS	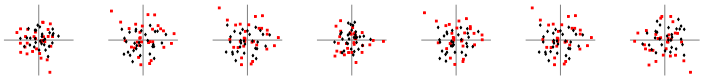						
II	#PCs	18	17	14	20	26	14	19
	MNV	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	t-test	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	PCS							
III	#PCs	19	18	15	28	30	14	20
	MNV	<1e-08	<1e-08	<1e-08	0.001	<1e-08	<1e-08	0.001
	t-test	<1e-08	3e-05	1e-07	8e-08	<1e-08	<1e-08	4e-06
	PCS	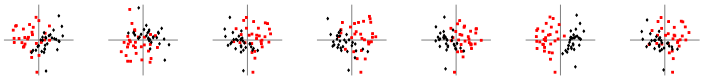						

Table S2.4 Model-independent comparison for model size 200, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) hypothesis tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replaced with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	24	32	25	36	43	26	39
	MNV	0.528	0.258	0.548	0.105	0.746	0.577	0.704
	t-test	0.530	0.836	0.804	0.784	0.378	0.805	0.879
	PCS							
II	#PCs	24	29	24	12	43	25	38
	MNV	<1e-08	<1e-08	5e-08	<1e-08	0.467	1e-07	<1e-08
	t-test	0.042	<1e-08	0.108	<1e-08	0.017	0.109	0.390
	PCS							
III	#PCs	17	7	8	13	38	1	31
	MNV	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	*	<1e-08
	t-test	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	PCS							

Table S2.5 Model-independent comparison for model size 400, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) hypothesis tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replaced with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	21	18	15	28	32	15	21
	MNV	0.740	0.936	0.849	0.776	0.854	0.850	0.990
	t-test	0.919	0.925	0.942	0.856	0.778	0.941	0.992
	PCS							
II	#PCs	10	10	8	8	15	8	11
	MNV	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	t-test	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	PCS							
III	#PCs	15	14	11	24	24	8	16
	MNV	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	t-test	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	PCS							

Table S2.6 Model-independent comparison for model size 400, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) hypothesis tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replaced with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	25	32	25	36	43	26	39
	MNV	0.521	0.817	0.617	0.004	0.338	0.630	0.276
	t-test	0.316	0.341	0.307	0.502	0.577	0.307	0.262
	PCS							
II	#PCs	23	24	24	1	43	25	37
	MNV	<1e-08	<1e-08	<1e-08	*	0.013	<1e-08	<1e-08
	t-test	<1e-08	<1e-08	<1e-08	<1e-08	1e-05	<1e-08	<1e-08
	PCS							
III	#PCs	9	1	1	1	26	1	21
	MNV	<1e-08	*	*	*	<1e-08	*	<1e-08
	t-test	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	PCS							

Table S2.7 Model-independent comparison for model size 800, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) hypothesis tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replaced with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	20	17	14	28	31	15	20
	MNV	0.242	0.274	0.252	0.005	0.379	0.285	0.322
	t-test	0.597	0.979	0.819	0.813	0.788	0.819	0.927
	PCS	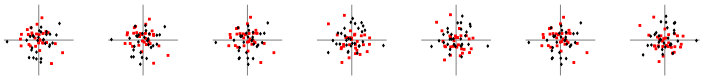						
II	#PCs	3	4	3	1	3	3	4
	MNV	<1e-08	<1e-08	<1e-08	*	<1e-08	<1e-08	<1e-08
	t-test	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	PCS							
III	#PCs	8	8	6	15	13	3	8
	MNV	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	t-test	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08	<1e-08
	PCS							

Table S2.8 Model-independent comparison for model size 800, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) hypothesis tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replaced with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	% var. (PC1)	15.6	10.0	12.4	9.0	8.2	12.3	8.8
	% var. (PC2)	9.1	9.2	9.7	6.6	6.5	9.6	8.0
	% var. (PC3)	7.4	7.3	7.8	5.6	4.9	7.8	6.5
	% var. (PC4)	6.7	6.2	6.8	5.3	4.2	6.7	5.6
	t-test (PC1)	0.036	0.057	0.049	0.023	0.503	0.049	0.077
	t-test (PC2)	0.680	0.532	0.495	0.550	0.358	0.495	0.317
	t-test (PC3)	0.997	0.762	0.853	0.941	0.445	0.849	0.863
	t-test (PC4)	0.714	0.716	0.904	0.546	0.090	0.910	0.958
	t-test* (PC1)	0.228	0.564	0.393	0.252	1.000	0.399	0.876
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
II	% var. (PC1)	17.5	11.7	13.9	19.6	8.5	13.8	10.0
	% var. (PC2)	9.2	8.9	9.8	7.4	6.0	9.7	8.1
	% var. (PC3)	7.9	7.0	8.2	5.6	4.5	8.1	6.6
	% var. (PC4)	7.3	6.8	7.6	5.3	4.2	7.6	6.3
	t-test (PC1)	0.001	0.001	0.003	$\leq$1e-08	0.956	0.003	0.003
	t-test (PC2)	0.766	0.886	0.884	0.829	0.096	0.885	0.942
	t-test (PC3)	0.441	0.874	0.355	0.728	0.967	0.357	0.536
	t-test (PC4)	0.515	0.250	0.287	0.887	0.160	0.285	0.220
	t-test* (PC1)	0.007	0.005	0.018	$\leq$1e-08	1.000	0.019	0.033
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
III	% var. (PC1)	16.1	25.4	24.4	22.6	7.1	55.0	11.9
	% var. (PC2)	9.8	7.8	8.9	6.4	6.0	5.1	8.7
	% var. (PC3)	8.4	6.6	7.5	6.0	5.3	4.2	7.5
	% var. (PC4)	8.1	6.5	7.3	5.6	4.8	4.1	7.1
	t-test (PC1)	6e-08	$\leq$1e-08	$\leq$1e-08	$\leq$1e-08	0.003	$\leq$1e-08	$\leq$1e-08
	t-test (PC2)	0.604	0.915	0.944	0.888	0.257	0.951	0.944
	t-test (PC3)	0.951	0.851	0.921	0.820	1e-06	0.972	0.983
	t-test (PC4)	0.510	0.842	0.802	0.893	0.312	0.892	0.704
	t-test* (PC1)	4e-07	$\leq$1e-08	$\leq$1e-08	$\leq$1e-08	0.041	$\leq$1e-08	$\leq$1e-08
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	2e-05	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000

Table S3.1 Percentage of explained variance (% var.) and t-test p-values before (t-test) and after (t-test*) weighted Bonferroni correction of the first four PCs for model size 100, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	% var. (PC1)	17.5	21.1	21.9	13.3	14.6	21.9	19.6
	% var. (PC2)	14.7	14.5	15.4	11.1	11.7	15.4	14.0
	% var. (PC3)	11.6	11.9	12.7	8.8	9.1	12.7	11.5
	% var. (PC4)	7.8	9.1	9.5	7.6	6.7	9.5	8.8
	t-test (PC1)	0.108	0.128	0.106	0.111	0.067	0.106	0.099
	t-test (PC2)	0.579	0.039	0.073	0.419	0.354	0.072	0.110
	t-test (PC3)	0.018	0.288	0.148	0.054	0.048	0.150	0.157
	t-test (PC4)	0.609	0.418	0.448	0.513	0.725	0.449	0.407
	t-test* (PC1)	0.616	0.608	0.484	0.837	0.461	0.486	0.505
	t-test* (PC2)	1.000	0.268	0.471	1.000	1.000	0.466	0.781
	t-test* (PC3)	0.158	1.000	1.000	0.611	0.529	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
II	% var. (PC1)	18.5	19.4	20.4	20.4	15.9	20.4	18.0
	% var. (PC2)	10.2	11.0	12.0	8.6	10.2	12.0	11.4
	% var. (PC3)	9.6	10.5	10.8	7.6	8.0	10.8	10.2
	% var. (PC4)	8.4	9.0	9.2	6.9	6.8	9.2	8.4
	t-test (PC1)	≤1e-08	1e-06	3e-08	≤1e-08	≤1e-08	3e-08	2e-07
	t-test (PC2)	0.010	0.786	0.087	0.588	0.245	0.087	0.181
	t-test (PC3)	0.597	0.004	0.190	0.095	0.669	0.190	0.402
	t-test (PC4)	0.299	0.089	0.060	0.667	0.034	0.061	0.045
	t-test* (PC1)	2e-08	5e-06	2e-07	≤1e-08	≤1e-08	2e-07	9e-07
	t-test* (PC2)	0.096	1.000	0.723	1.000	1.000	0.730	1.000
	t-test* (PC3)	1.000	0.034	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	0.990	0.656	1.000	0.500	0.661	0.538
III	% var. (PC1)	17.6	18.6	19.9	13.2	13.9	19.1	18.0
	% var. (PC2)	13.3	15.5	16.1	11.3	10.4	15.7	14.7
	% var. (PC3)	10.6	10.7	11.7	8.9	8.5	11.9	10.5
	% var. (PC4)	7.6	8.6	9.0	7.5	6.7	8.7	8.2
	t-test (PC1)	0.037	0.472	0.322	0.324	0.054	0.116	0.346
	t-test (PC2)	0.028	0.067	0.038	0.046	0.052	0.001	0.041
	t-test (PC3)	0.113	0.028	0.020	0.141	0.116	0.002	0.027
	t-test (PC4)	0.316	0.452	0.327	0.022	0.894	0.105	0.283
	t-test* (PC1)	0.211	1.000	1.000	1.000	0.385	0.608	1.000
	t-test* (PC2)	0.209	0.434	0.234	0.411	0.495	0.006	0.279
	t-test* (PC3)	1.000	0.257	0.171	1.000	1.000	0.013	0.254
	t-test* (PC4)	1.000	1.000	1.000	0.290	1.000	1.000	1.000

Table S3.2 Percentage of explained variance (% var.) and t-test p-values before (t-test) and after (t-test*) weighted Bonferroni correction of the first four PCs for model size 100, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	% var. (PC1)	18.2	12.1	14.5	11.7	6.7	14.3	10.3
	% var. (PC2)	8.5	8.3	9.0	6.6	5.5	9.0	7.3
	% var. (PC3)	7.3	7.4	8.0	5.5	4.7	7.9	6.5
	% var. (PC4)	6.6	6.7	6.9	5.4	4.4	6.9	5.9
	t-test (PC1)	0.783	0.810	0.895	0.973	0.830	0.895	0.979
	t-test (PC2)	0.505	0.391	0.460	0.304	0.151	0.459	0.332
	t-test (PC3)	0.141	0.198	0.150	0.856	0.400	0.150	0.248
	t-test (PC4)	0.546	0.553	0.541	0.427	0.234	0.536	0.409
	t-test* (PC1)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
II	% var. (PC1)	18.5	11.8	14.7	47.3	8.7	14.5	10.1
	% var. (PC2)	9.3	9.4	10.0	6.2	5.9	9.9	8.3
	% var. (PC3)	8.3	7.2	8.9	3.7	4.6	8.8	6.9
	% var. (PC4)	6.3	6.4	6.7	3.2	4.1	6.6	5.6
	t-test (PC1)	0.616	0.056	0.481	≤1e-08	0.773	0.480	0.534
	t-test (PC2)	0.440	0.077	0.489	0.463	0.915	0.490	0.433
	t-test (PC3)	0.476	0.075	0.448	0.994	0.617	0.448	0.488
	t-test (PC4)	0.036	2e-05	0.050	0.889	0.831	0.050	0.030
	t-test* (PC1)	1.000	0.473	1.000	≤1e-08	1.000	1.000	1.000
	t-test* (PC2)	1.000	0.820	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	0.576	4e-04	0.742	1.000	1.000	0.748	0.534
III	% var. (PC1)	17.1	49.8	43.9	43.8	14.5	80.7	13.4
	% var. (PC2)	10.6	6.2	7.4	5.0	6.0	2.4	9.8
	% var. (PC3)	9.8	5.2	6.6	4.9	5.1	2.3	7.9
	% var. (PC4)	9.1	4.1	5.0	4.2	4.2	1.8	6.4
	t-test (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test (PC2)	0.345	0.965	0.939	0.820	0.018	0.954	0.785
	t-test (PC3)	4e-05	0.800	0.574	0.559	0.008	0.819	0.245
	t-test (PC4)	0.291	0.544	0.317	0.545	0.418	0.731	0.764
	t-test* (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test* (PC2)	1.000	1.000	1.000	1.000	0.309	1.000	1.000
	t-test* (PC3)	5e-04	1.000	1.000	1.000	0.154	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000

Table S3.3 Percentage of explained variance (% var.) and t-test p-values before (t-test) and after (t-test*) weighted Bonferroni correction of the first four PCs for model size 200, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
	% var. (PC1)	17.3	20.2	21.0	13.5	13.7	20.9	18.8
	% var. (PC2)	16.2	16.7	18.2	12.6	12.0	18.2	16.5
	% var. (PC3)	9.7	9.4	10.0	9.2	7.8	10.1	9.0
	% var. (PC4)	8.6	8.5	9.2	6.4	7.0	9.1	8.3
	t-test (PC1)	0.192	0.294	0.280	0.985	0.378	0.280	0.240
I	t-test (PC2)	0.362	0.109	0.163	0.194	0.136	0.162	0.210
	t-test (PC3)	0.996	0.050	0.036	0.565	0.301	0.036	0.023
	t-test (PC4)	0.014	0.445	0.737	0.015	0.025	0.746	0.877
	t-test* (PC1)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC2)	1.000	0.650	0.891	1.000	1.000	0.892	1.000
	t-test* (PC3)	1.000	0.528	0.362	1.000	1.000	0.360	0.260
	t-test* (PC4)	0.166	1.000	1.000	0.238	0.360	1.000	1.000
	% var. (PC1)	32.7	30.6	34.1	48.1	33.6	34.0	29.8
	% var. (PC2)	9.0	10.9	11.1	7.7	7.0	11.1	10.8
	% var. (PC3)	8.7	9.4	9.0	5.8	6.6	9.0	8.7
	% var. (PC4)	7.1	6.9	7.3	3.9	5.7	7.3	6.6
	t-test (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
II	t-test (PC2)	0.703	0.430	0.553	0.798	0.563	0.552	0.589
	t-test (PC3)	0.887	0.703	0.451	0.772	0.664	0.450	0.306
	t-test (PC4)	0.584	0.840	0.936	0.674	0.933	0.937	0.857
	t-test* (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	% var. (PC1)	21.4	20.1	22.4	15.9	17.6	29.0	20.4
	% var. (PC2)	14.9	18.0	18.6	12.7	12.4	16.2	17.0
	% var. (PC3)	9.8	9.9	9.9	10.1	7.2	9.5	9.0
	% var. (PC4)	7.4	7.8	8.3	6.2	6.6	7.3	7.6
	t-test (PC1)	≤1e-08	3e-05	1e-07	8e-08	≤1e-08	≤1e-08	4e-06
III	t-test (PC2)	0.171	0.003	0.060	0.039	0.416	0.502	0.044
	t-test (PC3)	0.609	0.092	0.111	0.950	0.066	0.027	0.141
	t-test (PC4)	0.611	0.229	0.742	0.087	0.448	0.381	0.651
	t-test* (PC1)	≤1e-08	2e-04	6e-07	5e-07	≤1e-08	≤1e-08	2e-05
	t-test* (PC2)	1.000	0.019	0.323	0.304	1.000	1.000	0.261
	t-test* (PC3)	1.000	0.929	1.000	1.000	0.919	0.286	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000

Table S3.4 Percentage of explained variance (% var.) and t-test p-values before (t-test) and after (t-test*) weighted Bonferroni correction of the first four PCs for model size 200, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
	% var. (PC1)	19.1	13.0	15.9	11.7	8.2	15.8	10.9
	% var. (PC2)	8.6	7.5	8.7	6.4	6.2	8.6	6.4
	% var. (PC3)	7.1	6.8	7.5	5.6	4.5	7.4	6.1
	% var. (PC4)	6.4	6.5	6.6	5.3	4.3	6.6	5.8
	t-test (PC1)	0.530	0.836	0.804	0.784	0.378	0.805	0.879
I	t-test (PC2)	0.021	0.057	0.012	0.115	0.915	0.012	0.041
	t-test (PC3)	0.319	0.608	0.491	0.081	0.372	0.496	0.389
	t-test (PC4)	0.763	0.200	0.872	0.473	0.327	0.868	0.524
	t-test* (PC1)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC2)	0.240	0.757	0.135	1.000	1.000	0.136	0.633
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	% var. (PC1)	20.1	19.3	16.4	77.2	8.0	16.2	11.4
	% var. (PC2)	10.7	10.5	11.1	2.6	6.1	11.0	8.8
	% var. (PC3)	7.0	8.0	7.3	1.9	4.7	7.2	6.2
	% var. (PC4)	5.9	5.3	6.3	1.3	4.1	6.3	5.5
	t-test (PC1)	0.042	≤1e-08	0.108	≤1e-08	0.017	0.109	0.390
II	t-test (PC2)	0.128	0.029	0.106	0.720	0.175	0.107	0.086
	t-test (PC3)	0.004	0.110	0.006	0.885	0.909	0.006	0.002
	t-test (PC4)	0.767	0.401	0.966	0.978	0.060	0.954	0.631
	t-test* (PC1)	0.209	≤1e-08	0.661	≤1e-08	0.211	0.669	1.000
	t-test* (PC2)	1.000	0.276	0.955	1.000	1.000	0.971	0.975
	t-test* (PC3)	0.055	1.000	0.079	1.000	1.000	0.077	0.028
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	% var. (PC1)	38.6	79.2	74.3	73.6	32.4	94.0	29.1
	% var. (PC2)	9.4	2.6	3.4	2.8	4.4	0.8	7.9
	% var. (PC3)	7.6	2.4	3.2	2.4	4.2	0.7	7.7
	% var. (PC4)	6.9	2.0	2.7	2.0	3.7	0.6	5.7
	t-test (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
III	t-test (PC2)	0.183	0.966	0.792	0.662	0.979	0.878	0.792
	t-test (PC3)	0.734	0.906	0.826	0.976	0.777	0.951	0.754
	t-test (PC4)	0.602	0.800	0.587	0.827	0.979	0.875	0.555
	t-test* (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000

Table S3.5 Percentage of explained variance (% var.) and t-test p-values before (t-test) and after (t-test*) weighted Bonferroni correction of the first four PCs for model size 400, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	% var. (PC1)	16.4	19.2	20.0	12.7	13.1	19.9	17.8
	% var. (PC2)	13.9	16.5	17.0	11.5	11.2	16.9	15.1
	% var. (PC3)	12.1	10.8	11.8	9.8	8.0	11.8	11.0
	% var. (PC4)	8.8	9.2	9.9	8.5	6.8	9.8	8.9
	t-test (PC1)	0.919	0.925	0.942	0.856	0.778	0.941	0.992
	t-test (PC2)	0.694	0.474	0.492	0.867	0.444	0.490	0.483
	t-test (PC3)	0.278	0.228	0.268	0.514	0.250	0.268	0.298
	t-test (PC4)	0.923	0.884	0.998	0.350	0.721	0.999	0.956
	t-test* (PC1)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
II	% var. (PC1)	66.1	62.7	66.7	80.2	68.2	66.6	61.2
	% var. (PC2)	5.1	6.8	6.3	2.3	3.9	6.3	6.6
	% var. (PC3)	4.1	5.0	4.7	1.9	3.0	4.7	4.9
	% var. (PC4)	3.2	4.2	3.9	1.7	2.5	3.9	4.2
	t-test (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test (PC2)	0.726	0.631	0.631	0.908	0.739	0.631	0.608
	t-test (PC3)	0.601	0.763	0.773	0.838	0.827	0.774	0.842
	t-test (PC4)	0.732	0.694	0.651	0.963	0.799	0.650	0.600
	t-test* (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
III	% var. (PC1)	49.1	46.8	51.3	36.3	43.7	65.5	45.1
	% var. (PC2)	8.3	10.2	9.9	9.7	7.2	6.8	10.1
	% var. (PC3)	6.1	7.4	7.2	7.0	5.3	5.1	7.1
	% var. (PC4)	5.7	5.7	5.6	5.5	4.3	3.8	5.7
	t-test (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test (PC2)	0.833	0.855	0.963	0.572	0.854	0.970	0.968
	t-test (PC3)	0.686	0.522	0.527	0.923	0.510	0.560	0.519
	t-test (PC4)	0.709	0.711	0.814	0.623	0.985	0.815	0.749
	t-test* (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000

Table S3.6 Percentage of explained variance (% var.) and t-test p-values before (t-test) and after (t-test*) weighted Bonferroni correction of the first four PCs for model size 400, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
	% var. (PC1)	16.4	10.6	13.0	10.4	6.8	12.9	9.1
	% var. (PC2)	9.0	8.6	9.6	6.4	5.5	9.5	7.7
	% var. (PC3)	7.8	7.7	8.2	5.8	4.9	8.2	6.9
	% var. (PC4)	6.4	6.0	6.8	5.2	4.5	6.7	5.5
	t-test (PC1)	0.316	0.341	0.307	0.502	0.577	0.307	0.262
I	t-test (PC2)	0.376	0.320	0.414	0.146	0.823	0.414	0.359
	t-test (PC3)	0.783	0.785	0.832	0.613	0.280	0.837	0.872
	t-test (PC4)	0.496	0.646	0.491	0.141	0.581	0.492	0.507
	t-test* (PC1)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	% var. (PC1)	23.3	44.2	18.2	93.1	9.6	18.0	12.0
	% var. (PC2)	8.7	5.9	9.1	0.6	5.0	9.0	7.8
	% var. (PC3)	6.3	4.1	7.0	0.5	4.5	6.9	5.7
	% var. (PC4)	5.9	3.9	6.6	0.4	4.2	6.6	5.6
	t-test (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	1e-05	≤1e-08	≤1e-08
II	t-test (PC2)	0.005	0.609	0.015	0.906	0.066	0.015	0.060
	t-test (PC3)	0.819	0.792	0.956	0.948	0.395	0.952	0.759
	t-test (PC4)	0.775	0.913	0.826	0.959	0.485	0.827	0.875
	t-test* (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	1e-04	≤1e-08	≤1e-08
	t-test* (PC2)	0.055	1.000	0.162	1.000	1.000	0.163	0.767
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	% var. (PC1)	73.6	94.0	92.6	92.6	64.2	98.6	60.0
	% var. (PC2)	3.7	0.8	1.0	0.7	2.5	0.2	4.6
	% var. (PC3)	2.8	0.6	0.8	0.6	2.3	0.1	3.6
	% var. (PC4)	2.7	0.5	0.6	0.5	1.9	0.1	2.6
	t-test (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
III	t-test (PC2)	0.813	0.966	0.934	0.920	0.697	0.972	0.789
	t-test (PC3)	0.723	0.996	0.993	0.880	0.830	0.994	0.954
	t-test (PC4)	0.819	0.981	0.804	0.986	0.876	0.933	0.891
	t-test* (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000

Table S3.7 Percentage of explained variance (% var.) and t-test p-values before (t-test) and after (t-test*) weighted Bonferroni correction of the first four PCs for model size 800, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
	% var. (PC1)	18.0	21.2	21.8	13.9	14.3	21.8	19.8
	% var. (PC2)	16.4	19.0	19.6	12.3	13.6	19.6	17.7
	% var. (PC3)	11.0	10.4	11.4	9.1	9.1	11.4	10.8
	% var. (PC4)	7.9	8.6	8.9	6.6	6.8	8.9	8.2
	t-test (PC1)	0.597	0.979	0.819	0.813	0.788	0.819	0.927
I	t-test (PC2)	0.236	0.168	0.187	0.150	0.150	0.187	0.181
	t-test (PC3)	0.035	0.028	0.026	0.038	0.071	0.026	0.021
	t-test (PC4)	0.170	0.295	0.227	0.144	0.139	0.230	0.215
	t-test* (PC1)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC2)	1.000	0.885	0.953	1.000	1.000	0.954	1.000
	t-test* (PC3)	0.319	0.266	0.224	0.418	0.776	0.225	0.198
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	% var. (PC1)	87.4	85.1	87.4	94.0	88.3	87.3	84.2
	% var. (PC2)	1.8	2.3	2.1	0.7	1.2	2.1	2.4
	% var. (PC3)	1.4	2.0	1.8	0.6	1.1	1.8	2.0
	% var. (PC4)	1.3	1.8	1.6	0.5	1.0	1.6	1.8
	t-test (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
II	t-test (PC2)	0.987	0.952	0.978	0.996	0.999	0.978	0.988
	t-test (PC3)	0.945	0.917	0.963	0.973	0.932	0.963	0.955
	t-test (PC4)	0.975	0.996	0.988	0.986	0.972	0.988	0.995
	t-test* (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	% var. (PC1)	75.9	72.8	76.4	66.0	71.7	86.0	71.8
	% var. (PC2)	4.2	5.6	5.1	4.2	3.7	2.8	5.5
	% var. (PC3)	2.9	3.7	3.3	3.9	2.8	1.8	3.7
	% var. (PC4)	2.1	2.7	2.5	2.7	2.0	1.7	2.7
	t-test (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
III	t-test (PC2)	0.889	0.758	0.810	0.884	0.813	0.836	0.790
	t-test (PC3)	0.755	0.810	0.857	0.748	0.794	0.788	0.891
	t-test (PC4)	0.713	0.791	0.772	0.962	0.595	0.511	0.840
	t-test* (PC1)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	t-test* (PC2)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC3)	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	t-test* (PC4)	1.000	1.000	1.000	1.000	1.000	1.000	1.000

Table S3.8 Percentage of explained variance (% var.) and t-test p-values before (t-test) and after (t-test*) weighted Bonferroni correction of the first four PCs for model size 800, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs (30% var.)	3	4	4	5	6	4	5
	#PCs (50% var.)	7	8	7	11	13	7	9
	#PCs (70% var.)	12	15	13	19	24	13	18
	#PCs (90% var.)	26	32	26	37	43	27	38
	MNV (30% var.)	0.211	0.383	0.365	0.322	0.553	0.365	0.517
	MNV (50% var.)	0.652	0.159	0.664	0.507	0.545	0.663	0.202
	MNV (70% var.)	0.215	0.312	0.193	0.729	0.663	0.192	0.406
	MNV (90% var.)	0.852	0.509	0.851	0.966	0.285	0.879	0.486
II	#PCs (30% var.)	3	4	3	3	6	3	4
	#PCs (50% var.)	6	8	6	8	13	6	9
	#PCs (70% var.)	12	15	12	16	24	12	18
	#PCs (90% var.)	25	32	25	34	43	26	38
	MNV (30% var.)	0.011	0.010	0.019	≤1e-08	0.273	0.019	0.032
	MNV (50% var.)	0.041	0.060	0.044	≤1e-08	0.176	0.044	0.169
	MNV (70% var.)	0.070	0.013	0.152	≤1e-08	0.405	0.150	0.262
	MNV (90% var.)	0.023	4e-05	0.276	≤1e-08	0.448	0.256	0.496
III	#PCs (30% var.)	3	2	2	3	6	1	4
	#PCs (50% var.)	6	5	5	7	13	1	7
	#PCs (70% var.)	10	11	9	14	24	5	15
	#PCs (90% var.)	22	26	21	32	42	15	35
	MNV (30% var.)	2e-06	≤1e-08	≤1e-08	≤1e-08	2e-08	*	3e-08
	MNV (50% var.)	9e-06	≤1e-08	≤1e-08	≤1e-08	≤1e-08	*	1e-06
	MNV (70% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	1e-08	≤1e-08	5e-05
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	0.001	≤1e-08	6e-06

Table S4.1 Number of principal components (#PCs) and MANOVA test p-values (MNV) for several percentages of explained variance for model size 100, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replace with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs (30% var.)	2	2	2	3	3	2	2
	#PCs (50% var.)	4	4	3	6	6	4	4
	#PCs (70% var.)	8	7	6	13	13	6	8
	#PCs (90% var.)	20	19	15	29	31	15	22
	MNV (30% var.)	0.236	0.035	0.051	0.071	0.040	0.051	0.068
	MNV (50% var.)	0.065	0.073	0.043	0.150	0.198	0.068	0.085
	MNV (70% var.)	0.292	0.303	0.189	0.106	0.570	0.189	0.425
	MNV (90% var.)	0.116	0.141	0.175	0.085	0.121	0.175	0.125
II	#PCs (30% var.)	3	2	2	3	3	2	3
	#PCs (50% var.)	5	5	4	6	6	4	5
	#PCs (70% var.)	9	8	7	11	12	7	8
	#PCs (90% var.)	20	19	15	27	29	15	22
	MNV (30% var.)	≤1e-08	7e-06	2e-08	≤1e-08	≤1e-08	2e-08	8e-07
	MNV (50% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	MNV (70% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	4e-08
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	5e-06
III	#PCs (30% var.)	2	2	2	3	3	2	2
	#PCs (50% var.)	5	4	4	6	6	4	4
	#PCs (70% var.)	9	8	7	13	13	7	8
	#PCs (90% var.)	21	18	15	29	30	15	22
	MNV (30% var.)	0.009	0.145	0.069	0.064	0.015	0.001	0.078
	MNV (50% var.)	0.020	0.050	0.016	0.034	0.036	4e-06	0.021
	MNV (70% var.)	0.090	0.106	0.042	0.270	0.001	2e-08	0.078
	MNV (90% var.)	0.006	0.001	0.110	0.139	2e-04	≤1e-08	0.147

Table S4.2 Number of principal components (#PCs) and MANOVA test p-values (MNV) for several percentages of explained variance for model size 100, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replace with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs (30% var.)	3	4	3	5	7	3	4
	#PCs (50% var.)	6	8	6	10	14	7	9
	#PCs (70% var.)	12	15	12	19	25	12	18
	#PCs (90% var.)	25	32	25	36	43	26	38
	MNV (30% var.)	0.446	0.598	0.456	0.863	0.117	0.456	0.570
	MNV (50% var.)	0.788	0.102	0.767	0.293	0.521	0.846	0.143
	MNV (70% var.)	0.142	0.208	0.158	0.800	0.762	0.159	0.522
	MNV (90% var.)	0.287	0.214	0.157	0.123	0.699	0.187	0.356
II	#PCs (30% var.)	3	4	3	1	6	3	4
	#PCs (50% var.)	6	8	6	2	13	6	9
	#PCs (70% var.)	11	15	12	8	24	12	18
	#PCs (90% var.)	24	31	25	27	43	25	38
	MNV (30% var.)	0.719	2e-06	0.674	*	0.847	0.674	0.186
	MNV (50% var.)	0.400	≤1e-08	0.471	≤1e-08	0.780	0.471	0.547
	MNV (70% var.)	0.618	≤1e-08	0.800	≤1e-08	0.305	0.799	0.457
	MNV (90% var.)	0.022	≤1e-08	0.276	≤1e-08	0.778	0.270	0.038
III	#PCs (30% var.)	3	1	1	1	5	1	3
	#PCs (50% var.)	5	2	2	3	11	1	7
	#PCs (70% var.)	9	6	6	9	21	1	15
	#PCs (90% var.)	21	19	17	27	41	6	35
	MNV (30% var.)	≤1e-08	*	*	*	≤1e-08	*	≤1e-08
	MNV (50% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	*	6e-08
	MNV (70% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	*	≤1e-08
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	3e-08

Table S4.3 Number of principal components (#PCs) and MANOVA test p-values (MNV) for several percentages of explained variance for model size 200, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replace with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs (30% var.)	2	2	2	3	3	2	2
	#PCs (50% var.)	4	4	4	6	6	4	4
	#PCs (70% var.)	9	8	7	13	14	7	8
	#PCs (90% var.)	21	19	15	29	31	15	21
	MNV (30% var.)	0.282	0.158	0.209	0.575	0.252	0.209	0.227
	MNV (50% var.)	0.072	0.082	0.102	0.205	0.165	0.102	0.084
	MNV (70% var.)	0.121	0.106	0.273	0.317	0.274	0.279	0.099
	MNV (90% var.)	0.138	0.264	0.322	0.342	0.652	0.322	0.244
II	#PCs (30% var.)	1	1	1	1	1	1	2
	#PCs (50% var.)	3	3	3	2	4	3	4
	#PCs (70% var.)	7	7	6	6	9	6	7
	#PCs (90% var.)	18	17	14	20	26	14	19
	MNV (30% var.)	*	*	*	*	*	*	≤1e-08
	MNV (50% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	MNV (70% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
III	#PCs (30% var.)	2	2	2	3	2	2	2
	#PCs (50% var.)	4	4	3	5	6	3	4
	#PCs (70% var.)	8	7	6	11	12	6	8
	#PCs (90% var.)	19	18	15	28	30	14	20
	MNV (30% var.)	≤1e-08	6e-07	6e-08	1e-07	≤1e-08	≤1e-08	1e-06
	MNV (50% var.)	5e-08	4e-07	3e-08	1e-07	≤1e-08	≤1e-08	4e-06
	MNV (70% var.)	1e-06	≤1e-08	5e-07	3e-05	≤1e-08	≤1e-08	9e-07
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	0.001	≤1e-08	≤1e-08	0.001

Table S4.4 Number of principal components (#PCs) and MANOVA test p-values (MNV) for several percentages of explained variance for model size 200, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replace with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs (30% var.)	3	4	3	5	6	3	5
	#PCs (50% var.)	6	8	6	10	13	7	9
	#PCs (70% var.)	11	15	12	19	25	12	19
	#PCs (90% var.)	24	32	25	36	43	26	39
	MNV (30% var.)	0.080	0.235	0.078	0.185	0.696	0.078	0.297
	MNV (50% var.)	0.267	0.514	0.278	0.020	0.530	0.311	0.614
	MNV (70% var.)	0.241	0.611	0.344	0.044	0.682	0.344	0.527
	MNV (90% var.)	0.528	0.258	0.548	0.105	0.746	0.577	0.704
II	#PCs (30% var.)	2	3	3	1	6	3	4
	#PCs (50% var.)	6	6	6	1	13	6	9
	#PCs (70% var.)	11	12	11	1	24	12	18
	#PCs (90% var.)	24	29	24	12	43	25	38
	MNV (30% var.)	0.038	≤1e-08	0.004	*	0.035	0.004	0.007
	MNV (50% var.)	2e-06	≤1e-08	3e-04	*	0.034	3e-04	3e-04
	MNV (70% var.)	4e-05	≤1e-08	0.002	*	0.049	0.001	2e-06
	MNV (90% var.)	≤1e-08	≤1e-08	5e-08	≤1e-08	0.467	1e-07	≤1e-08
III	#PCs (30% var.)	1	1	1	1	1	1	2
	#PCs (50% var.)	3	1	1	1	6	1	4
	#PCs (70% var.)	6	1	1	1	17	1	11
	#PCs (90% var.)	17	7	8	13	38	1	31
	MNV (30% var.)	*	*	*	*	*	*	≤1e-08
	MNV (50% var.)	≤1e-08	*	*	*	≤1e-08	*	≤1e-08
	MNV (70% var.)	≤1e-08	*	*	*	≤1e-08	*	≤1e-08
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	*	≤1e-08

Table S4.5 Number of principal components (#PCs) and MANOVA test p-values (MNV) for several percentages of explained variance for model size 400, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replace with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs (30% var.)	2	2	2	3	3	2	2
	#PCs (50% var.)	4	4	4	6	7	4	4
	#PCs (70% var.)	8	7	6	12	14	6	8
	#PCs (90% var.)	21	18	15	28	32	15	21
	MNV (30% var.)	0.921	0.772	0.790	0.924	0.579	0.788	0.783
	MNV (50% var.)	0.858	0.743	0.796	0.962	0.486	0.795	0.818
	MNV (70% var.)	0.836	0.801	0.775	0.897	0.734	0.775	0.895
	MNV (90% var.)	0.740	0.936	0.849	0.776	0.854	0.850	0.990
II	#PCs (30% var.)	1	1	1	1	1	1	1
	#PCs (50% var.)	1	1	1	1	1	1	1
	#PCs (70% var.)	2	3	2	1	2	2	3
	#PCs (90% var.)	10	10	8	8	15	8	11
	MNV (30% var.)	*	*	*	*	*	*	*
	MNV (50% var.)	*	*	*	*	*	*	*
	MNV (70% var.)	≤1e-08	≤1e-08	≤1e-08	*	≤1e-08	≤1e-08	≤1e-08
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
III	#PCs (30% var.)	1	1	1	1	1	1	1
	#PCs (50% var.)	2	2	1	3	2	1	2
	#PCs (70% var.)	5	4	4	8	7	2	5
	#PCs (90% var.)	15	14	11	24	24	8	16
	MNV (30% var.)	*	*	*	*	*	*	*
	MNV (50% var.)	≤1e-08	≤1e-08	*	≤1e-08	≤1e-08	*	≤1e-08
	MNV (70% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08

Table S4.6 Number of principal components (#PCs) and MANOVA test p-values (MNV) for several percentages of explained variance for model size 400, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replace with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs (30% var.)	3	4	3	5	7	3	5
	#PCs (50% var.)	6	8	7	10	14	7	10
	#PCs (70% var.)	12	15	12	19	25	13	19
	#PCs (90% var.)	25	32	25	36	43	26	39
	MNV (30% var.)	0.605	0.709	0.629	0.373	0.496	0.629	0.765
	MNV (50% var.)	0.879	0.633	0.904	0.295	0.265	0.904	0.738
	MNV (70% var.)	0.506	0.530	0.757	0.014	0.483	0.795	0.687
	MNV (90% var.)	0.521	0.817	0.617	0.004	0.338	0.630	0.276
II	#PCs (30% var.)	2	1	3	1	6	3	4
	#PCs (50% var.)	6	2	6	1	13	6	9
	#PCs (70% var.)	11	8	12	1	24	12	17
	#PCs (90% var.)	23	24	24	1	43	25	37
	MNV (30% var.)	<1e-08	*	<1e-08	*	1e-06	<1e-08	<1e-08
	MNV (50% var.)	<1e-08	<1e-08	<1e-08	*	2e-04	<1e-08	<1e-08
	MNV (70% var.)	<1e-08	<1e-08	<1e-08	*	0.001	<1e-08	<1e-08
	MNV (90% var.)	<1e-08	<1e-08	<1e-08	*	0.013	<1e-08	<1e-08
III	#PCs (30% var.)	1	1	1	1	1	1	1
	#PCs (50% var.)	1	1	1	1	1	1	1
	#PCs (70% var.)	1	1	1	1	4	1	4
	#PCs (90% var.)	9	1	1	1	26	1	21
	MNV (30% var.)	*	*	*	*	*	*	*
	MNV (50% var.)	*	*	*	*	*	*	*
	MNV (70% var.)	*	*	*	*	<1e-08	*	<1e-08
	MNV (90% var.)	<1e-08	*	*	*	<1e-08	*	<1e-08

Table S4.7 Number of principal components (#PCs) and MANOVA test p-values (MNV) for several percentages of explained variance for model size 800, parameter set 1, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replace with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs (30% var.)	2	2	2	3	3	2	2
	#PCs (50% var.)	4	3	3	6	6	3	4
	#PCs (70% var.)	8	7	6	13	13	6	8
	#PCs (90% var.)	20	17	14	28	31	15	20
	MNV (30% var.)	0.433	0.389	0.410	0.090	0.142	0.410	0.410
	MNV (50% var.)	0.086	0.078	0.078	0.087	0.237	0.078	0.067
	MNV (70% var.)	0.158	0.240	0.182	0.022	0.363	0.182	0.205
	MNV (90% var.)	0.242	0.274	0.252	0.005	0.379	0.285	0.322
II	#PCs (30% var.)	1	1	1	1	1	1	1
	#PCs (50% var.)	1	1	1	1	1	1	1
	#PCs (70% var.)	1	1	1	1	1	1	1
	#PCs (90% var.)	3	4	3	1	3	3	4
	MNV (30% var.)	*	*	*	*	*	*	*
	MNV (50% var.)	*	*	*	*	*	*	*
	MNV (70% var.)	*	*	*	*	*	*	*
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	*	≤1e-08	≤1e-08	≤1e-08
III	#PCs (30% var.)	1	1	1	1	1	1	1
	#PCs (50% var.)	1	1	1	1	1	1	1
	#PCs (70% var.)	1	1	1	2	1	1	1
	#PCs (90% var.)	8	8	6	15	13	3	8
	MNV (30% var.)	*	*	*	*	*	*	*
	MNV (50% var.)	*	*	*	*	*	*	*
	MNV (70% var.)	*	*	*	≤1e-08	*	*	*
	MNV (90% var.)	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08

Table S4.8 Number of principal components (#PCs) and MANOVA test p-values (MNV) for several percentages of explained variance for model size 800, parameter set 2, and $n = 30$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replace with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.841	0.340	0.293	0.499	0.471	0.828
	arg max	0.593	0.273	0.301	0.170	0.596	0.518
	min	0.598	0.896	0.900	0.326	0.012	0.356
	arg min	0.284	0.343	0.519	1.000	0.555	1.000
	$\overline{X}^{ss}$	0.995	0.415	0.967	0.679	0.017	0.962
	S^{ss}	0.402	0.326	0.415	0.593	0.343	0.415
II	max	0.014	3e−4	0.024	0.023	0.494	0.171
	arg max	0.019	0.307	0.077	0.206	0.140	0.025
	min	0.063	0.173	0.209	0.237	0.017	0.037
	arg min	0.720	0.184	0.017	1.000	0.294	1.000
	$\overline{X}^{ss}$	0.244	$\leq$ 1e−8	0.012	$\leq$ 1e−8	0.456	0.013
	S^{ss}	0.751	0.685	0.826	0.196	0.404	0.826
III	max	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	1e−7	4e−4	$\leq$ 1e−8
	arg max	0.066	0.345	0.001	0.009	0.044	0.015
	min	0.004	8e−5	5e−5	0.949	0.010	$\leq$ 1e−8
	arg min	0.047	0.402	0.015	1.000	0.044	2e−5
	$\overline{X}^{ss}$	8e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8
	S^{ss}	0.010	0.014	0.074	0.001	0.388	0.457

Table S5.1 P-values for the empirical selection of FMs for model size 400, parameter set 1, $n = 10$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{ss}$ and S^{ss} statistics, and with the Mann-Whitney U test for the arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	12	14	13	15	16	13	15
	MNV	0.591	0.120	0.503	0.858	0.927	0.505	0.790
	t-test	0.662	0.575	0.539	0.492	0.364	0.539	0.504
	PCS							
II	#PCs	12	13	12	7	16	12	14
	MNV	0.046	2e-05	0.118	≤1e-08	0.410	0.118	0.056
	t-test	0.365	8e-07	0.723	≤1e-08	0.039	0.725	0.902
	PCS	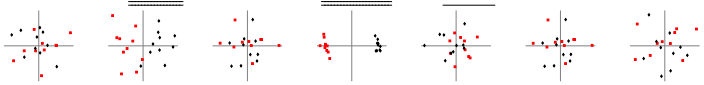						
III	#PCs	10	5	6	7	14	1	13
	MNV	≤1e-08	≤1e-08	≤1e-08	≤1e-08	0.001	*	5e-04
	t-test	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	1e-08
	PCS							

Table S5.2 Results for the model-independent comparison for model size 400, parameter set 1, and $n = 10$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) statistical tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined. The MANOVA test is not applicable to one PC, and the respective p-values are replace with an asterisk (*).

Comp.	Stat.	Outputs					
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$
I	max	0.289	0.085	0.093	0.249	0.205	0.595
	arg max	0.024	0.159	0.018	0.090	0.779	0.535
	min	0.991	0.406	0.679	0.622	0.918	0.264
	arg min	0.152	0.445	0.177	1.000	0.005	0.158
	$\overline{X}^{ss}$	0.547	0.052	0.491	0.080	0.209	0.481
	S^{ss}	0.259	0.100	0.228	0.973	0.467	0.228
II	max	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	0.042	3e−7
	arg max	$\leq$ 1e−8	0.143	0.078	3e−4	0.155	$\leq$ 1e−8
	min	$\leq$ 1e−8	$\leq$ 1e−8	9e−7	0.225	0.056	$\leq$ 1e−8
	arg min	0.324	0.225	$\leq$ 1e−8	1.000	0.136	0.158
	$\overline{X}^{ss}$	3e−5	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	0.034	$\leq$ 1e−8
	S^{ss}	0.082	0.097	0.136	2e−4	0.862	0.138
III	max	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8
	arg max	$\leq$ 1e−8	5e−5	$\leq$ 1e−8	0.015	0.024	$\leq$ 1e−8
	min	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	0.248	0.001	$\leq$ 1e−8
	arg min	$\leq$ 1e−8	4e−7	$\leq$ 1e−8	1.000	8e−5	$\leq$ 1e−8
	$\overline{X}^{ss}$	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8
	S^{ss}	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	$\leq$ 1e−8	2e−5

Table S6.1 P-values for the empirical selection of FMs for model size 400, parameter set 1, and $n = 100$ runs per configuration. P-values were obtained with the t-test for max, min, $\overline{X}^{ss}$ and S^{ss} statistics, and with the Mann-Whitney U test for the arg max and arg min statistics. Case I weights two similar configurations, case II compares configurations with a small implementation difference, and case III compares configurations with a different parameter. Values lower than 0.05 are underlined, while values lower than 0.01 are double-underlined.

Comp.	Data	Outputs						
		P^s	P^w	P^c	$\overline{E}^s$	$\overline{E}^w$	$\overline{C}$	$\widetilde{A}$
I	#PCs	35	56	36	72	108	38	93
	MNV	0.750	0.164	0.756	1e-05	0.521	0.657	3e-05
	t-test	0.522	0.528	0.479	0.164	0.909	0.479	0.352
	PCS	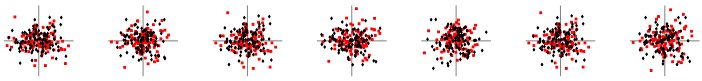						
II	#PCs	35	49	36	19	107	37	91
	MNV	≤1e-08	≤1e-08	≤1e-08	≤1e-08	3e-06	≤1e-08	≤1e-08
	t-test	≤1e-08	≤1e-08	≤1e-08	≤1e-08	1e-06	≤1e-08	2e-08
	PCS							
III	#PCs	22	10	11	20	86	1	66
	MNV	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	*	≤1e-08
	t-test	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08	≤1e-08
	PCS	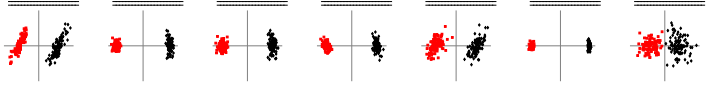						

Table S6.2 Results for the model-independent comparison for model size 400, parameter set 1, and $n = 100$ runs per configuration. Case I weights two similar setups, case II compares setups with a small implementation difference, and case III compares setups with a different parameter. #PCs is the number of PCs which explain at least 90% of variance, while MNV and t-test refer to the p-value yielded by the MANOVA (#PCs) and t (first PC) statistical tests, respectively. PCS shows the representation of the respective output in the first two dimensions of the PC space. Output $\widetilde{A}$ refers to the concatenation of all outputs (range scaled). The MANOVA test is not applicable to one PC, and the respective p-values are replace with an asterisk (*). P-values lower than 0.05 are underlined, while p-values lower than 0.01 are double-underlined.